

Combinatorial Inference for Graphical Models

Matey Neykov^{*†}

Junwei Lu^{*†}

Han Liu^{*‡}

Abstract

We propose a new family of combinatorial inference problems for graphical models. Unlike classical statistical inference where the main interest is point estimation or parameter testing, combinatorial inference aims at testing the global structure of the underlying graph. Examples include testing the graph connectivity, the presence of a cycle of certain size, or the maximum degree of the graph. To begin with, we develop a unified theory for the fundamental limits of a large family of combinatorial inference problems. We propose new concepts including structural packing and buffer entropies to characterize how the complexity of combinatorial graph structures impacts the corresponding minimax lower bounds. On the other hand, we propose a family of novel and practical structural testing algorithms to match the lower bounds. We provide thorough numerical results on both synthetic graphical models and brain networks to illustrate the usefulness of these proposed methods.

Keywords: Graph structural inference; minimax testing; uncertainty assessment; multiple hypothesis testing; post-regularization inference.

1 Introduction

Graphical model is a powerful tool for modeling complex relationships among many random variables. A central theme of graphical model research is to infer the structure of the underlying graph based on observational data. Though significant progress has been made,

^{*}Department of Operations Research and Financial Engineering, Princeton University, Princeton NJ 08544, USA; e-mail: {mneykov,junwei,hanliu}@princeton.edu

[†]These authors contributed equally to this work.

[‡]Research supported by NSF DMS1454377-CAREER; NSF IIS 1546482-BIGDATA; NIH R01MH102339; NSF IIS1408910; NIH R01GM083084.

existing works mainly focus on estimating the graph (Meinshausen and Bühlmann, 2006; Liu et al., 2009; Ravikumar et al., 2011; Cai et al., 2011) or testing the existence of a single edge (Jankova and van de Geer, 2015; Ren et al., 2015; Neykov et al., 2015; Gu et al., 2015).

In this paper we consider a new inferential problem: testing the combinatorial structure of the underlying graph. Examples include testing the graph connectivity, cycle presence, or assessing the maximum degree of the graph. Unlike classical inference which aims at testing a set of Euclidean parameters, combinatorial inference aims to test some global structural properties and requires the development of new methodology. As for methodological development, this paper mainly considers the Gaussian graphical model (though our method is applicable to the more general semiparametric exponential family graphical models and elliptical copula graphical models): Let $\mathbf{X} = (X_1, \dots, X_d)^T \sim N_d(\mathbf{0}, (\boldsymbol{\Theta}^*)^{-1})$ be a d -dimensional Gaussian random vector with precision matrix $\boldsymbol{\Theta}^* = (\Theta_{jk}^*)$. Let $G^* = G(\boldsymbol{\Theta}^*) := (\bar{V}, E^*)$ be an undirected graph, where $\bar{V} = \{1, \dots, d\}$ and an edge $(j, k) \in E^*$ if and only if $\Theta_{jk}^* \neq 0$. It is well known that G^* has the pairwise Markov property, i.e., $(j, k) \notin E^*$ if and only if X_j and X_k are conditionally independent given the remaining variables. In a combinatorial inference problem, our goal is to test whether G^* has certain global structural properties based on n random samples $\mathbf{X}_1, \dots, \mathbf{X}_n$. Specifically, let \mathcal{G} be the set of all graphs over the vertex set \bar{V} and $(\mathcal{G}_0, \mathcal{G}_1)$ be a pair of non-overlapping subsets of \mathcal{G} . We assume all the graphs in \mathcal{G}_1 have a property (e.g., connectivity) while the graphs in \mathcal{G}_0 do not have this property. Such a pair $(\mathcal{G}_0, \mathcal{G}_1)$ is called a sub-decomposition of \mathcal{G} . Our goal is to test the hypothesis $\mathbf{H}_0 : G^* \in \mathcal{G}_0$ versus $\mathbf{H}_1 : G^* \in \mathcal{G}_1$. We provide several concrete examples below.

Connectivity. A graph is connected if and only if there exists a path connecting each pair of its vertices. To test connectivity, we set $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ is disconnected}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ is connected}\}$. Under the Gaussian graphical model, this is equivalent to testing whether the variables can be partitioned into at least two independent sets.

Cycle presence. Sometimes it is of interest to test whether the underlying graph is acyclic. In this example we let $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ is cycle-free}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ contains a cycle}\}$. If a graph is acyclic, it can be easily visualized on a two dimensional plane.

Maximum degree. Another relevant question is to test whether the maximum degree of the graph is less than or equal to some integer $s_0 \in \mathbb{N}$ versus the alternative that the maximum degree is at least $s_1 \in \mathbb{N}$, where $s_0 < s_1$. Define the sub-decomposition $\mathcal{G}_0 = \{G : d_{\max}(G) \leq s_0\}$ and $\mathcal{G}_1 = \{G : d_{\max}(G) \geq s_1\}$ respectively.

While our ultimate goal is to test whether $G^* \in \mathcal{G}_0$ versus $G^* \in \mathcal{G}_1$, our access to G^* is only through the random samples $\{\mathbf{X}_i\}_{i=1}^n$. Under Gaussian models, we can translate the original problem of testing graphs to testing the precision matrix:

$$\mathbf{H}_0 : \boldsymbol{\Theta}^* \in \mathcal{S}_0 \text{ vs } \mathbf{H}_1 : \boldsymbol{\Theta}^* \in \mathcal{S}_1. \quad (1.1)$$

In (1.1), $\mathcal{S}_0, \mathcal{S}_1 \subset \mathcal{M}(s)$ are two sets of precision matrices such that for all $\Theta \in \mathcal{S}_0, \mathcal{S}_1$ we have $G(\Theta) \in \mathcal{G}_0, \mathcal{G}_1$ respectively, and $\mathcal{M}(s)$ is defined as

$$\mathcal{M}(s) = \{\Theta \in \mathbb{R}^{d \times d} : \Theta = \Theta^T, C^{-1} \leq \Theta \leq C, \|\Theta\|_1 \leq L, \max_{j \in [d]} \|\Theta_{*j}\|_0 \leq s\} \quad (1.2)$$

for some constants $1 \leq C \leq L$. The set $\mathcal{M}(s)$ restricts our attention to well conditioned symmetric matrices Θ , whose induced graphs $G(\Theta)$ have maximum degree of at most s . Given this setup, we aim to characterize necessary conditions on the pair $\mathcal{S}_0, \mathcal{S}_1$ under which the combinatorial inference problem in (1.1) is *testable*. Specifically, recall that a test is any measurable function $\psi : \{\mathbf{X}_i\}_{i=1}^n \mapsto \{0, 1\}$. Define the risk of testing \mathcal{S}_0 against \mathcal{S}_1 as:

$$\gamma(\mathcal{S}_0, \mathcal{S}_1) = \inf_{\psi} \left[\max_{\Theta \in \mathcal{S}_0} \mathbb{P}_{\Theta}(\psi = 1) + \max_{\Theta \in \mathcal{S}_1} \mathbb{P}_{\Theta}(\psi = 0) \right]. \quad (1.3)$$

If $\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0, \mathcal{S}_1) = 1$, we say that the problem (1.1) is *untestable* since any test fails to distinguish between \mathcal{S}_0 and \mathcal{S}_1 in the asymptotic minimax sense. Due to the close relationship between the sets of precision matrices $(\mathcal{S}_0, \mathcal{S}_1)$ and the sub-decomposition $(\mathcal{G}_0, \mathcal{G}_1)$ (recall that for all $\Theta \in \mathcal{S}_0, \mathcal{S}_1$ we have $G(\Theta) \in \mathcal{G}_0, \mathcal{G}_1$ resp.), we anticipate that the sub-decomposition can capture the intrinsic challenge of the test in (1.1). Indeed, in Sections 2 and 4 we develop a new theory to characterize the impact of the combinatorial structures of \mathcal{G}_0 and \mathcal{G}_1 to the lower bound of the risk $\gamma(\mathcal{S}_0, \mathcal{S}_1)$. Such lower bounds provide necessary conditions for any valid test. We then develop practical procedures that match the obtained lower bounds. To understand how the sub-decomposition $(\mathcal{G}_0, \mathcal{G}_1)$ affects the intrinsic difficulty of the problem in (1.1), we consider the three examples given before. Our theory on lower bounds distinguishes between two types of sub-decompositions — in the first type one can find graphs belonging to \mathcal{G}_0 and \mathcal{G}_1 differing in only one single edge, while in the second type all graphs belonging to \mathcal{G}_0 must differ on multiple edge sets from the graphs belonging to \mathcal{G}_1 .

One can check that in the first two examples (connectivity and cycle presence testing) there exist graphs belonging to \mathcal{G}_0 and \mathcal{G}_1 differing in only one single edge. For instance, when testing connectivity, consider a tree with a single edge removed (thus it becomes a forest) versus a connected tree. Extending this intuition, for a fixed graph $G_0 = (\bar{V}, E_0) \in \mathcal{G}_0$, we call the edge set $\mathcal{C} = \{e_1, \dots, e_m\}$ a single-edge *null-alternative separator* (NAS) if for all edges $e \in \mathcal{C}$ the graphs $(\bar{V}, E_0 \cup \{e\}) \in \mathcal{G}_1$. Intuitively the bigger the cardinality of a NAS set is, the harder it is to tell the null from the alternatives. In Section 2 we detail that $\gamma(\mathcal{S}_0, \mathcal{S}_1)$ is asymptotically 1, when the signal strength of separation between \mathcal{S}_0 and \mathcal{S}_1 is low (see (2.2) and (2.3) for a formal definition) and there exists a NAS set with sufficiently large packing number. The packing number, formally defined in Definition 2.3, represents the cardinality of a subset of edges in \mathcal{C} which are “far” apart, where the proximity measure of two edges is a predistance¹ based on the graph G_0 . Recall that for a graph G and two

¹Compared to distance, a predistance does not have to satisfy the triangle inequality.

vertices u and v , a graph geodesic distance is defined by:

$$d_G(u, v) := \text{length of the shortest path between } u \text{ and } v \text{ within } G.$$

Using the notion of geodesic distance, one can define a predistance between two edges, by taking the minimum over the geodesic distances of their corresponding nodes.

If the difference between null and alternative is more than one edge, as in the maximum degree testing $\mathcal{G}_0 = \{G : d_{\max}(G) \leq s_0\}$ vs $\mathcal{G}_1 = \{G : d_{\max}(G) \geq s_1\}$ for example, the packing number does not always capture the lower bound of the tests. In Section 4, we develop a novel mechanism to handle this more sophisticated case. We introduce a new concept called “buffer entropy” which can produce sharper lower bounds to overcome the disadvantages of the packing number.

On the other hand, to match the lower bound, we propose the null-alternative witness technique (NAWT) as a unified approach for combinatorial tests. The NAWT identifies a critical structure and proceeds to test whether this structure indeed belongs to the true graph. We prove the NAWT tests can control both the type I and type II errors asymptotically.

1.1 Contributions

There are three major contributions of this paper.

Our first contribution is to develop a novel strategy for obtaining minimax lower bounds on the signal strength required to distinguish combinatorial graph structures which are separable via a single-edge NAS. In particular, our theory bridges the information-theoretic lower bounds to the packing number of the NAS set, which is an intuitive combinatorial quantity. To obtain this result, we relate the chi-square divergence between two probability measures to the number of “closed walks” on their corresponding Markov graphs. Our analysis hinges on several technical tools including Le Cam’s Lemma, matrix perturbation inequalities and spectral graph theory. The usefulness of the approach is demonstrated by obtaining generic and interpretable lower bounds in numerous examples such as testing connectivity, connected components, self-avoiding paths, and cycles.

Our second contribution is to provide a device for proving lower bounds under the settings where the null and alternative graphs differ in multiple edges. Under such case, the packing number does not always provide tight lower bounds. Our technical device is based on a new graph combinatorial quantity called buffer entropy. The buffer entropy is a complexity measure of the structural tests and characterizes their lower bounds. We apply this device to derive lower bounds for testing the maximum degree and detecting a sparse clique.

Our third contribution is to propose a null-alternative witness technique (NAWT) for testing (1.1), which matches the lower bounds on the signal strength. The NAWT is a two step procedure — in the first step it identifies a minimal structure witnessing the alternative hypothesis, and in the second step it attempts to certify the presence of this structure in the

graph. NAWT utilizes recent advances in high-dimensional inference and provides honest tests for combinatorial inference problems. It has two advantages compared to the support recovery procedures in [Meinshausen and Bühlmann \(2006\)](#); [Ravikumar et al. \(2011\)](#); [Cai et al. \(2011\)](#): First, it allows us to control the type-I error at any given level; Second, it does not require us to perfectly recover the underlying graph to conduct valid inference.

1.2 Related Work

Graphical model inference is relatively straightforward when $d < n$, but becomes notoriously challenging when $d \gg n$. In high-dimensions, estimation procedures were studied by [Yuan and Lin \(2006\)](#); [Friedman et al. \(2008\)](#); [Lam and Fan \(2009\)](#); [Cai et al. \(2011\)](#) among others, while for variable selection procedures see [Meinshausen and Bühlmann \(2006\)](#); [Raskutti et al. \(2008\)](#); [Liu et al. \(2009\)](#); [Ravikumar et al. \(2011\)](#); [Cai et al. \(2011\)](#) and references therein. Recently, motivated by [Zhang and Zhang \(2014\)](#), various inferential methods for high-dimensional graphical models were suggested ([Liu, 2013](#); [Jankova and van de Geer, 2015](#); [Chen et al., 2015](#); [Ren et al., 2015](#); [Neykov et al., 2015](#); [Gu et al., 2015](#), e.g.), most of which focus on testing the presence of a single edge (except [Liu \(2013\)](#) who took the FDR approach ([Benjamini and Hochberg, 1995](#)) to conduct multiple tests and [Gu et al. \(2015\)](#) who developed new procedures of edge testing in Gaussian copula models). None of the aforementioned works address the problem of combinatorial structure testing.

In addition to estimation and model selection procedures, efforts have been made to understand the fundamental limits of these problems. Lower bounds on estimation were obtained by [Ren et al. \(2015\)](#), where the authors show that the parametric estimation rate $n^{-1/2}$ is unattainable unless $s \log d / \sqrt{n} = o(1)$. Lower bounds on the minimal sample size required for model selection in Ising models were established by [Santhanam and Wainwright \(2012\)](#), where it is shown that support recovery is unattainable when $n \ll s^2 \log d$. In a follow up work, [Wang et al. \(2010\)](#) studied model selection limits on the sample size in Gaussian graphical models. The latter two works are remotely related to ours, in that both works exploit graph properties to obtain information-theoretic lower bounds. However, our problem differs significantly from theirs since we focus on developing lower bounds for testing graph structure, which is a fundamentally different problem from estimating the whole graph.

Our problem is most closely related to those in [Addario-Berry et al. \(2010\)](#); [Arias-Castro et al. \(2012, 2015a,b\)](#), which are inspired by the large body of research on minimax hypothesis testing ([Ingster, 1982](#); [Ingster et al., 2010](#); [Arias-Castro et al., 2011a,b](#), e.g.) among many others. [Addario-Berry et al. \(2010\)](#) quantify the signal strength as the mean parameter of a standard Gaussian distribution, while [Arias-Castro et al. \(2012, 2015a\)](#) impose models on the covariance matrix of a multivariate Gaussian distribution. In our setup the parameter spaces of interest are designed to reflect the graphical model structure, and hence the signal strength is naturally imposed on the precision matrix. [Arias-Castro et al. \(2015b\)](#) address

testing on a lattice based Gaussian Markov random field. For specific problems they establish lower bounds on the signal strength required to test the empty graph versus an alternative hypothesis. This is different from the setting of our problems, where the null hypothesis is usually not the empty graph.

1.3 Notation

The following notations are used throughout the paper. For a vector $\mathbf{v} = (v_1, \dots, v_d)^T \in \mathbb{R}^d$, let $\|\mathbf{v}\|_q = (\sum_{i=1}^d v_i^q)^{1/q}$, $1 \leq q < \infty$, $\|\mathbf{v}\|_0 = |\text{supp}(\mathbf{v})|$, where $\text{supp}(\mathbf{v}) = \{j : v_j \neq 0\}$, and $|A|$ denotes the cardinality of a set A . Furthermore, let $\|\mathbf{v}\|_\infty = \max_i |v_i|$ and $\mathbf{v}^{\otimes 2} = \mathbf{v}\mathbf{v}^T$. For a matrix \mathbf{A} , we denote \mathbf{A}_{*j} and \mathbf{A}_{j*} to be the j^{th} column and row of \mathbf{A} respectively. For any $n \in \mathbb{N}$ we use the shorthand notation $[n] = \{1, \dots, n\}$. For two integer sets $S_1, S_2 \subseteq [d]$, we denote $\mathbf{A}_{S_1 S_2}$ to be the sub-matrix of \mathbf{A} with elements $\{A_{jk}\}_{j \in S_1, k \in S_2}$. Moreover, we denote $\|\mathbf{A}\|_{\max} = \max_{jk} |A_{jk}|$, $\|\mathbf{A}\|_p = \max_{\|\mathbf{v}\|_p=1} \|\mathbf{A}\mathbf{v}\|_p$ for $p \geq 1$.

For a graph G we use $V(G)$, $E(G)$ to refer to the vertex set, edge set of G respectively. We also denote $V(E)$ as the vertex set of the edge set E . We reserve special notations for the complete vertex set $\bar{V} := [d]$, the complete edge set $\bar{E} := \{e \in 2^{[d]} \mid |e| = 2\}$ and the complete graph $\bar{G} := (\bar{V}, \bar{E})$. For two integers $j, k \in [d]$ we use unordered pairs $(j, k) = (k, j)$ to denote undirected edges between vertex k and vertex j . Any symmetric matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ naturally induces an undirected graph $G(\mathbf{A})$, with vertices in the set \bar{V} and edge set $E(G(\mathbf{A})) = \{(j, k) : A_{jk} \neq 0, j \neq k\}$. Additionally if E is an arbitrary edge set (i.e., $E \subseteq \bar{E}$) for $e := (j, k) \in E$ we use the notations $A_e = A_{jk} = A_{kj}$ interchangeably to denote the element e of the matrix \mathbf{A} .

Given two sequences $\{a_n\}, \{b_n\}$ we write $a_n = O(b_n)$ if there exists a constant $C < \infty$ such that $a_n \leq Cb_n$; $a_n = o(b_n)$ if $a_n/b_n \rightarrow 0$, and $a_n \asymp b_n$ if there exists positive constants c and C such that $c < a_n/b_n < C$. Finally we use the shorthand notations \wedge and \vee for min and max of two numbers respectively.

1.4 Organization of the Paper

The paper is structured as follows. Our main result along with applications to several examples is presented in Section 2. In Section 3 we outline the NAWT method, and illustrate how to apply our procedure to the examples considered in Section 2. In Section 4 we generalize the lower bounds strategies from the single edge NAS setting to multiple edge NAS setting. Thorough numerical studies and a real data analysis are presented in Section 5. A brief discussion is provided in Section 6. All proofs are deferred to the Appendix.

2 Single-Edge Null-Alternative Separators

In this section we derive a novel and generic lower bound strategy, applicable to null and alternative hypotheses which differ in one single edge: i.e., under the Gaussian model, there exist two matrices $\Theta_0 \in \mathcal{S}_0$ and $\Theta_1 \in \mathcal{S}_1$ whose induced graphs $G_0 := G(\Theta_0)$ and $G_1 := G(\Theta_1)$ differ in a single edge. We formalize this concept in the definition below.

Definition 2.1 (Single-Edge Null-Alternative Separator). For a sub-decomposition $(\mathcal{G}_0, \mathcal{G}_1)$ of \mathcal{G} , let $G_0 = (\bar{V}, E_0) \in \mathcal{G}_0$ be a graph under the null. We refer to an edge set

$$\mathcal{C} = \{e_1, \dots, e_m\},$$

as a (single-edge) *null-alternative separator* (NAS) with the null base G_0 if for any $e \in \mathcal{C}$ the graphs $G_e := (\bar{V}, E_0 \cup \{e\}) \in \mathcal{G}_1$.

As remarked in the introduction, if a large NAS set exists it is expected that differentiating G_0 from an alternative graph $G_e \in \mathcal{G}_1$ is more challenging. Indeed, our main result of this section confirms this intuition. We proceed to define a predistance for a graph G and two edges e, e' (which need not belong to $E(G)$) which plays a key role in our lower bound result.

Definition 2.2 (Edge Geodesic Predistance). Let $G = (\bar{V}, E)$ and $\{e, e'\}$ be a pair of edges (e and e' may or may not belong to E). We define

$$d_G(e, e') := \min_{u \in e, v \in e'} d_G(u, v),$$

where $d_G(u, v)$ denotes the geodesic distance between vertices u and v on the augmented graph G . If such a path does not exist $d_G(e, e') = \infty$.

$d_G(e, e')$ is a predistance, i.e., $d_G(e, e) = 0$ and $d_G(e, e') \geq 0$. Moreover, $d_G(e, e')$ has the same value regardless of whether $e, e' \in E(G)$. See Figure 1 for an illustration of $d_G(e, e')$.

To sharply characterize combinatorial inference lower bounds, we need to define a notion of structural packing entropy for graphs. The classical packing entropy on metric spaces has been widely used in statistics community (Yang and Barron, 1999, e.g.) to characterize the information-theoretic lower bound for parametric estimation. We propose the structural packing entropy for graphs as an analog of the classical packing entropy on metric spaces.

Definition 2.3 (Structural Packing Entropy). Let \mathcal{C} be a non-empty edge set and G be a graph. For any $r \geq 0$ we call the edge set $N_r \subset \mathcal{C}$ an r -packing of \mathcal{C} if for any $e, e' \in N_r$ we have $d_G(e, e') \geq r$. Define the structural r -packing entropy as:

$$M(\mathcal{C}, d_G, r) := \log \max \{|N_r| \mid N_r \subset \mathcal{C}, N_r \text{ is a } r\text{-packing of } \mathcal{C}\}. \quad (2.1)$$

The packing entropy in Definition 2.3 is an analog to the classical packing entropy on metric spaces in the sense that it is defined over an edge set \mathcal{C} equipped with a predistance $d_G(e, e')$ based on the graph G . Similar to the relation between the packing entropy in metric spaces and the lower bounds for parametric estimation (Yang and Barron, 1999), we will show that the structural packing entropy plays a pivotal role in characterizing the information-theoretic lower bound for the graph structural tests.

To study minimax lower bounds, we only need to focus on the Gaussian graphical model whose structural properties are completely characterized by the precision matrices. We now formally define the sets of precision matrices \mathcal{S}_0 and \mathcal{S}_1 used in this section. Let:

$$\mathcal{S}_0(\theta, s) := \left\{ \Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_0, \min_{e \in E(G(\Theta))} |\Theta_e| \geq \theta \right\} \text{ and,} \quad (2.2)$$

$$\mathcal{S}_1(\theta, s) := \left\{ \Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_1, \min_{e \in E(G(\Theta))} |\Theta_e| \geq \theta \right\}, \quad (2.3)$$

where $\mathcal{M}(s)$ is defined in (1.2). The parameter θ in the definitions of $\mathcal{S}_0(\theta, s)$ and $\mathcal{S}_1(\theta, s)$ denotes the *signal strength*, and as we show below, its magnitude plays an important role in determining whether one can distinguish between graphical models in $\mathcal{S}_0(\theta, s)$ and $\mathcal{S}_1(\theta, s)$.

Theorem 2.1 (Necessary Signal Strength). Let $G_0 \in \mathcal{G}_0$ be a graph whose maximum degree is bounded by an integer D , and let \mathcal{C} be a single-edge NAS with null base G_0 . Suppose that $n \geq (\log |\mathcal{C}|)/c_0$ for some $c_0 > 0$, and

$$\theta \leq \kappa \sqrt{\frac{M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|)}{n}} \wedge \frac{1 - C^{-1}}{\sqrt{2}(D + 2)}, \quad (2.4)$$

where C is defined in (1.2). Then if $M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|) \rightarrow \infty$ as $n \rightarrow \infty$, there exists a sufficiently small constant κ in (2.4) (depending on D, C, L, c_0) such that:

$$\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1.$$

Theorem 2.1 parallels the classical minimax results in Yang and Barron (1999), but unlike their work we quantify the signal strength necessary for testing combinatorial structure rather than estimating Euclidean parameters. The radius $\log |\mathcal{C}|$ of the packing entropy in (2.4) ensures that the pairs of distinct edges are sufficiently far apart. The constant term $\frac{1 - C^{-1}}{\sqrt{2}(D + 2)}$ in (2.4) ensures that precision matrices with signal strength θ indeed belong to $\mathcal{M}(s)$.

Proof Sketch. The proof of Theorem 2.1 can roughly be divided into four steps. More details will be provided in Appendix A.

Step 1 (*Connect the structural parameters to geometric parameters*). Given the adjacency matrices of the null and alternative graphs G_0 and $\{G_e\}_{e \in \mathcal{C}}$, we construct the corresponding precision matrices and make sure that they belong to $\mathcal{S}_0(\theta, s)$ and $\mathcal{S}_1(\theta, s)$.

Step 2 (*Construct a risk lower bound using Le Cam's method*). The second step uses Le Cam's method to lower bound the risk $\gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s))$. This requires us to evaluate the chi-square divergence between a normal and a mixture normal distribution. The chi-square divergence can be expressed via ratios of determinants. In particular, we show that the log chi-square divergence can be equivalently re-expressed via an infinite sum of differences among trace operators of adjacency matrix powers.

Step 3 (*Represent the lower bound by the number of shortest closed walks in the graph*). In this step we control the deviations of the differences of the trace operators. Since the trace of the power of an adjacency matrix equals the number of closed walks within the corresponding graph, we eliminate the trace powers which are smaller than the shortest closed walks. The traces of the higher powers are handled via matrix perturbation bounds.

Step 4 (*Characterize the smallest magnitude of the geometric parameter using the packing entropy*). Lastly, we show that condition (2.4) ensures that the closed walks on the packing of the NAS set are sufficiently lengthy, which shows that the chi-square divergence vanishes when the signal strength θ is small. \square

2.1 Some Applications

In this section we give several examples of combinatorial testing, which readily fall into the framework developed in Section 2.

Example 2.1 (Connectivity Testing). Consider the sub-decomposition $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ disconnected}\}$ vs $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ connected}\}$. We construct a base graph $G_0 := (\bar{V}, E_0)$ where

$$E_0 := \{(j, j+1)_{j=1}^{\lfloor d/2 \rfloor - 1}, (\lfloor d/2 \rfloor, 1), (j, j+1)_{j=\lfloor d/2 \rfloor + 1}^d, (\lfloor d/2 \rfloor + 1, d)\},$$

and let

$$\mathcal{C} := \{(j, \lfloor d/2 \rfloor + j)_{j=1}^{\lfloor d/2 \rfloor}\} \text{ (see Figure 1)}.$$

Clearly adding any edge from \mathcal{C} to G_0 connects the graph, so \mathcal{C} is a single edge NAS with a null base G_0 . Furthermore, the maximum degree of G_0 equals 2 by construction. To construct a packing set of \mathcal{C} , we collect all edges $(j, \lfloor d/2 \rfloor + j)$ satisfying $\lceil \log |\mathcal{C}| \rceil$ divides j except if $j > \lfloor d/2 \rfloor - \lceil \log |\mathcal{C}| \rceil$. This procedure results in a packing set with radius at least $\lceil \log |\mathcal{C}| \rceil$ which has cardinality of at least $\left\lfloor \frac{|\mathcal{C}|}{\lceil \log |\mathcal{C}| \rceil} \right\rfloor - 1$. Therefore

$$M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|) \geq \log \left[\left\lfloor \frac{|\mathcal{C}|}{\lceil \log |\mathcal{C}| \rceil} \right\rfloor - 1 \right] \asymp \log |\mathcal{C}| \asymp \log d.$$

Theorem 2.1 implies that testing connectivity is impossible when $\theta < \kappa \sqrt{\log d/n} \wedge \frac{1-C^{-1}}{4\sqrt{2}}$.




Figure 1: The graph G_0 with two edges $e, e' \in \mathcal{C} : d_{G_0}(e, e') = 2, d = 10$.

Example 2.2 ($\mathbf{m} + 1$ vs \mathbf{m} Connected Components, $\mathbf{m} \geq \sqrt{d}$). Let $\mathbf{m} \geq \sqrt{d}$ be an integer. In this example we are interested in testing whether the graph contains $\mathbf{m} + 1$ connected components vs \mathbf{m} connected components. The reason to assume $\mathbf{m} \geq \sqrt{d}$ is to make sure there are sufficiently many edges for constructing a single edge NAS in order to obtain sharp bounds. The case when $\mathbf{m} < \sqrt{d}$ is treated in Example 2.7 via a different NAS construction (In fact, the case $\mathbf{m} < \sqrt{d}$ requires deleting edges from the alternative rather than adding edges to the null base. See Section 2.2 for more details). Formally we have the sub-decomposition $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ has } \geq \mathbf{m} + 1 \text{ connected components}\}$ vs $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ has } \leq \mathbf{m} \text{ connected components}\}$. Construct the null base graph $G_0 = (\bar{V}, E_0)$, where

$$E_0 := \{(j, j + 1)_{j=1}^{d-\mathbf{m}-1}\},$$

and we let

$$\mathcal{C} := \{(j, j + 1)_{j=d-\mathbf{m}}^{d-1}\} \text{ (see Figure 2).}$$

Adding an edge $e \in \mathcal{C}$ to G_0 converts the base graph G_0 into a graph with \mathbf{m} connected components and therefore \mathcal{C} is a single edge NAS with a null base G_0 . Additionally, the maximum degree of G_0 is 2 by construction. Note that the distance between any two edges in \mathcal{C} is 0 if and only if they share a common vertex, and ∞ in all other cases. This implies that we can construct a packing set by taking every other edge in the set \mathcal{C} . We conclude that $M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|) \asymp \log(|\mathcal{C}|/2) \asymp \log d$. Hence, by Theorem 2.1 the minimax risk goes to 1 when $\theta < \kappa \sqrt{\log d/n} \wedge \frac{1-C^{-1}}{4\sqrt{2}}$.




Figure 2: Null base graph G_0 with $d - \mathbf{m} - 1$ edges, NAS \mathcal{C} (dashed), $d_{G_0}(e, e') = \infty, d = 8$.

Example 2.3 (Cycle Detection). Consider testing whether the graph is acyclic vs the graph contains a cycle. Let $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ is cycle-free}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ contains a cycle}\}$. Define the null base graph $G_0 = (\bar{V}, E_0)$

$$E_0 := \{(j, j+1)_{j=1}^{d-1}\},$$

and let the edge set

$$\mathcal{C} := \{(j, j+2)_{j=1}^d\},$$

where the addition is taken modulo d (refer to Figure 3 for a visualization). By construction we have $G_0 \in \mathcal{G}_0$ and $|\mathcal{C}| = d$. Adding any edge from \mathcal{C} to G_0 results in a graph with a cycle, and hence the edge set \mathcal{C} is a single edge NAS with a null base G_0 . The maximum degree of G_0 equals 2, and is thus bounded. Moreover, there exists a $\log |\mathcal{C}|$ -packing set of \mathcal{C} of cardinality at least $\frac{|\mathcal{C}|-2}{\lceil \log |\mathcal{C}| \rceil + 2}$ which can be produced by collecting the edges $(j, j+2)$ for $j = k(\log |\mathcal{C}| + 2) + 1$ for $k = 0, 1, \dots$ and $j \leq d-2$. The last observation implies that $M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|) \asymp \log \frac{|\mathcal{C}|-2}{\lceil \log |\mathcal{C}| \rceil + 2} \asymp \log d$. Hence by Theorem 2.1 we conclude that the minimax risk goes to 1 when $\theta < \kappa \sqrt{\log d/n} \wedge \frac{1-C^{-1}}{4\sqrt{2}}$.




Figure 3: The graph G_0 with two (dashed) edges $e, e' \in \mathcal{C}$ such that $d_{G_0}(e, e') = 2$, $d = 7$.

Example 2.4 (Tree vs Connected Graph with Cycles). The construction in Example 2.3 also shows that we have the same signal strength limitation to test for cycles, even if we restrict to the subclass of connected graphs, i.e., the class of graphs under the null hypothesis is the class of all trees $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ is a tree}\}$, and the alternative is the class of all connected graphs which contain a cycle — $\mathcal{G}_1 \in \{G \in \mathcal{G} : G \text{ is connected but is not a tree}\}$.

Example 2.5 (Triangle-Free Graph). Consider testing whether the graph contains a triangle (i.e., 3-clique). More formally let the decomposition of \mathcal{G} be $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ is triangle-free}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : \exists \text{ 3-clique subgraph of } G\}$. It is clear that in this case we can reuse the set \mathcal{C} and its null base $\tilde{G}_0 = (\bar{V}, E_0 \cup \{(1, d)\})$, where E_0 and \mathcal{C} are taken as in Example 2.3.

Example 2.6 (Self-Avoiding Path (SAP) of Length \mathfrak{m} vs $\mathfrak{m} + 1$, $\mathfrak{m} < \sqrt{d}$). In this example we test the existence of a SAP of length $\leq \mathfrak{m}$ vs SAP of length $\geq \mathfrak{m} + 1$, i.e., $\mathcal{G}_0 = \{G \in \mathcal{G} : \text{All SAPs have length} \leq \mathfrak{m}\}$ vs $\mathcal{G}_1 = \{G \in \mathcal{G} : \exists \text{ SAP of length } \mathfrak{m} + 1\}$. We assume that $\mathfrak{m} < \sqrt{d}$ in order to be able to construct sufficiently large NAS set which yields tight bounds. The case $\mathfrak{m} \geq \sqrt{d}$ is treated in Example 2.8. Without loss of generality we further suppose that $\mathfrak{m} + 2$ divides d (at the sake of assigning some isolated vertices in the graph). Construct the following graph in \mathcal{G}_0 : $G_0 = (\bar{V}, E_0)$, where

$$E_0 := \{(j, j + 1) : \text{unless } (\mathfrak{m} + 2) \text{ divides } j \text{ or } (\mathfrak{m} + 2) \text{ divides } (j + 1)\},$$

and hence $|E_0| = d\mathfrak{m}/(\mathfrak{m} + 2)$ (see Figure 4). Consider the class:

$$\mathcal{C} := \{(j(\mathfrak{m} + 2) - 1, j(\mathfrak{m} + 2))_{j \in [d/(\mathfrak{m} + 2)]}\}.$$

The set \mathcal{C} is a NAS since adding any edge from \mathcal{C} to G_0 results in a graph with a longer self-avoiding path. Furthermore the maximum degree of G_0 is 2. Thus when $\theta < \frac{1-C^{-1}}{4\sqrt{2}}$ we can apply Theorem 2.1. Notice that the set \mathcal{C} is its own a $\log |\mathcal{C}|$ -packing, since each two edges $e, e' \in \mathcal{C}$ satisfy $d_{G_0}(e, e') = \infty$. Hence $M(|\mathcal{C}|, d_{G_0}, \log |\mathcal{C}|) = \log |\mathcal{C}| \asymp \log d$. We conclude that if $\theta \leq \kappa \sqrt{\log d/n}$ we cannot differentiate the null from the alternative hypothesis.




Figure 4: The null base G_0 and the NAS \mathcal{C} (dashed) with $d_{G_0}(e, e') = \infty$, $d = 15$, $\mathfrak{m} = 3$.

2.2 Edge Deletion NAS Extension

Though Theorem 2.1 delivers valid lower bounds, the obtained bounds may not always be tight. For example, recall the construction of Example 2.6. When $\mathfrak{m} = d - 2$ the NAS set contains only a single edge, and therefore (2.4) is not sharp. Sometimes tighter bounds can be obtained via constructions which delete edges from the alternative graph instead of adding edges to the graph under null. In this subsection we extend the NAS definition to allow for deleting edges from the alternative hypothesis, which is sometimes more convenient and gives more informative lower bounds as we demonstrate below.

Definition 2.4 (Single-Edge Deletion Null-Alternative Separator). Let $G_1 = (\bar{V}, E_1)$ be an alternative graph. An edge set \mathcal{C} is called a (single edge) deletion NAS with alternative base G_1 , if for all $e \in \mathcal{C}$ the graphs $G_{\setminus e} = (\bar{V}, E_1 \setminus \{e\}) \in \mathcal{G}_0$.

In parallel to Theorem 2.1 we have the following result using the deletion NAS.

Theorem 2.2. Let $G_1 \in \mathcal{G}_1$ be a graph whose maximum degree bounded by a fixed integer D , and let \mathcal{C} be a single-edge deletion NAS with alternative base G_1 . Suppose that $n \geq (\log |\mathcal{C}|)/c_0$ for some $c_0 > 0$ and that:

$$\theta \leq \kappa \sqrt{\frac{M(\mathcal{C}, d_{G_1}, \log |\mathcal{C}|)}{n}} \wedge \frac{1 - C^{-1}}{\sqrt{2D}}. \quad (2.5)$$

Then, if $M(\mathcal{C}, d_{G_1}, \log |\mathcal{C}|) \rightarrow \infty$ as $n \rightarrow \infty$, for a sufficiently small κ (depending on D, C, L, c_0) we have:

$$\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1.$$

To illustrate the usefulness of Theorem 2.2 for edge deletion NAS, we consider the examples below.

Example 2.7 ($\mathfrak{m} + 1$ vs \mathfrak{m} Connected Components, $\mathfrak{m} < \sqrt{d}$). In this example we consider a complementary setting to Example 2.2, and test whether the graph contains $\mathfrak{m} + 1$ connected components vs \mathfrak{m} connected components for $\mathfrak{m} < \sqrt{d}$. Recall the decomposition $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ has } \geq \mathfrak{m} + 1 \text{ connected components}\}$ vs $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ has } \leq \mathfrak{m} \text{ connected components}\}$. Construct an alternative base graph $G_1 = (\bar{V}, E_1)$:

$$E_1 := \{(j, j + 1)_{j=1}^{d-\mathfrak{m}}\},$$

and set

$$\mathcal{C} := E_1 = \{(j, j + 1)_{j=1}^{d-\mathfrak{m}}\} \text{ (see Figure 5 for a visualization).}$$

Since removing any edge from \mathcal{C} results in increasing the number of connected components by 1, the set \mathcal{C} is a single edge deletion NAS with an alternative base graph G_1 . Moreover, notice that the maximum degree of G_1 is 2. The entropy $M(\mathcal{C}, d_{G_1}, \log |\mathcal{C}|)$ satisfies $M(\mathcal{C}, d_{G_1}, \log |\mathcal{C}|) \asymp \log \frac{|\mathcal{C}|}{\log |\mathcal{C}|} \asymp \log |\mathcal{C}| \asymp \log d$. Hence by Theorem 2.2 we conclude that testing connectivity is impossible when $\theta < \kappa \sqrt{\log d/n} \wedge \frac{1-C^{-1}}{2\sqrt{2}}$.



Figure 5: The graph G_1 with $d - \mathfrak{m}$ edges and $\mathfrak{m} - 1$ isolated nodes, $d = 8$.

Example 2.8 (Self-Avoiding Path (SAP) of Length \mathfrak{m} vs $\mathfrak{m} + 1$, $\mathfrak{m} \geq \sqrt{d}$). In parallel to Example 2.6 we now consider testing that $\text{SAP} \leq \mathfrak{m}$ vs SAP of length $\geq \mathfrak{m} + 1$ when $\mathfrak{m} \geq \sqrt{d}$. Define $\mathcal{G}_0 = \{G \in \mathcal{G} : \forall \text{ SAP have length } \leq \mathfrak{m}\}$ vs $\mathcal{G}_1 = \{G \in \mathcal{G} : \exists \text{ SAP of length } \mathfrak{m} + 1\}$. Construct the alternative graph $G_1 = (\bar{V}, E_1)$:

$$E_1 := \{(j, j + 1)_{j=1}^{\mathfrak{m}+1}\}.$$

Define the set

$$\mathcal{C} := E_1 = \{(j, j+1)_{j=1}^{m+1}\},$$

and note that removing any edge from \mathcal{C} results in a graph from the null. Hence \mathcal{C} is a deletion NAS with base G_1 . Moreover, the maximum degree of G_1 is 2. Similarly to Example 2.7 it follows that $M(\mathcal{C}, d_{G_1}, \log |\mathcal{C}|) \asymp \log |\mathcal{C}| \asymp \log d$. An application of Theorem 2.2 implies that for values of $\theta < \kappa \sqrt{\log d/n} \wedge \frac{1-C^{-1}}{2\sqrt{2}}$ testing the length of SAP is impossible.



Figure 6: The graph G_1 with $m+1$ edges and two edges $e, e' \in \mathcal{C}$ with $d_{G_1}(e, e') = 1$.

Remark 2.1. When $\Theta \in \mathcal{M}(s)$, the results in Theorem 2.1 and Theorem 2.2 suggest that a signal strength of order $\sqrt{\log d/n}$ is necessary for controlling the risk (1.3). In fact, Theorem 7 of Cai et al. (2011) shows that under such signal strength condition, support recovery of Θ is indeed achievable, which further implies that controlling the risk (1.3) is possible. A naive procedure for matching the lower bound is to first perfectly recover the graph structure. Then construct a test based on examining whether the graph has the desired combinatorial structure. Though such an approach is theoretically feasible, it is not practical. First, such an approach is overly conservative and does not allow us to tightly control the type-I error at a desired level. Second, such an approach crucially depends on having a suitable thresholding parameter to estimate the graph, which is in general not realistic.

In the next section, we present a family of novel test procedures, which do not require perfect support recovery of the full graph. Compared to the naive approach described in Remark 2.1, our tests explicitly exploit the combinatorial structure of the targeted hypotheses and can control the type-I error at any desired level.

3 Null-Alternative Witness Technique

We start with a high-level outline of a new combinatorial inference approach for graphical models. For clarity we mainly present using the case of Gaussian graphical models, and comment on extensions to other graphical models in Section 3.3.

We assume henceforth that the full graph $\bar{G} = (\bar{V}, \bar{E})$ belongs to \mathcal{G}_1 . For a graph $G \in \mathcal{G}_1$, define the following class of edge sets

$$\mathcal{W}_1(G) = \{E : \forall E', E \subseteq E' \subseteq E(G) \Rightarrow (\bar{V}, E') \in \mathcal{G}_1\}.$$

The set $\mathcal{W}_1(G)$ collects all edge sets forming graphs in \mathcal{G}_1 , which can be obtained by the graph G via iteratively pruning one edge at a time and maintaining that in each step the

resulting graph still belongs to \mathcal{G}_1 . We use the shorthand notation

$$\overline{\mathcal{W}}_1 = \mathcal{W}_1(\overline{G}).$$

Consider the following parameter sets:

$$\mathcal{S}_0(s) := \left\{ \Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_0 \right\}, \text{ and} \quad (3.1)$$

$$\mathcal{S}_1(\theta, s) := \left\{ \Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_1, \max_{E \in \overline{\mathcal{W}}_1} \min_{e \in E} |\Theta_e| \geq \theta \right\}, \quad (3.2)$$

The parameter space (3.1) does not impose any assumption on the minimum signal strength, thus is broader than the one defined in (2.2). In Definition (3.2), the signal strength is not imposed on all the edges of the alternative graphs. In fact, we only need to impose the signal strength assumption on a subset of edges which can be obtained by pruning the complete graph. Such a condition is much weaker than the usual condition needed for perfect graph recovery. We note that for sub-decompositions $(\mathcal{G}_0, \mathcal{G}_1)$ satisfying $\mathcal{W}_1(G) \subseteq \overline{\mathcal{W}}_1$ for all $G \in \mathcal{G}_1$, the parameter space (3.2) is strictly larger than the parameter space (2.3).

Given n independent samples $\mathbf{X}_i \sim N(0, (\Theta^*)^{-1})$ and a sub-decomposition $(\mathcal{G}_0, \mathcal{G}_1)$ we formulate a procedure for testing $\mathbf{H}_0 : \Theta^* \in \mathcal{S}_0(s)$ vs $\mathbf{H}_1 : \Theta^* \in \mathcal{S}_1(\theta, s)$. Let $\hat{\Sigma} := n^{-1} \sum_{i=1}^n \mathbf{X}_i^{\otimes 2}$ be the empirical covariance matrix. Let $\hat{\Theta}$ be any estimator of the precision matrix Θ^* satisfying for some fixed constant $K > 0$:

$$\|\hat{\Theta} - \Theta^*\|_{\max} \leq K \sqrt{\log d/n}, \quad (3.3)$$

$$\|\hat{\Theta} - \Theta^*\|_1 \leq K s \sqrt{\log d/n}, \quad \|\hat{\Sigma} \hat{\Theta} - \mathbf{I}_d\|_{\max} \leq K \sqrt{\log d/n}, \quad (3.4)$$

with probability at least $1 - d^{-1}$ uniformly over the parameter space $\mathcal{M}(s)$ (recall definition (1.2)). An example of an estimator of Θ^* with this properties is the CLIME procedure introduced by Cai et al. (2011) (see also (5.2)). An overview of the *null-alternative witness technique* (NAWT) is sketched below:

- i. In the first step NAWT identifies a *minimal structure witnessing the alternative*;
- ii. In the second step NAWT attempts to *certify* that the minimal structure identified by the first step is indeed present in the graph.

Split the data $\mathcal{D} = \{\mathbf{X}_i\}_{i=1}^n$ in two approximately equal-sized sets $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}$, $\mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$ and obtain estimates $\hat{\Theta}^{(1)}, \hat{\Theta}^{(2)}$ on \mathcal{D}_1 and \mathcal{D}_2 correspondingly. For the first step, we exploit $\hat{\Theta}^{(1)}$ to solve the following max-min combinatorial optimization problem

$$\hat{E} := \operatorname{argmax}_{E \in \overline{\mathcal{W}}_1} \min_{e \in E} |\hat{\Theta}_e^{(1)}|, \quad (3.5)$$

where edge sets with the smallest cardinality are preferred, and further ties are broken arbitrarily. Program (3.5) aims at identifying the smallest edge set in $\overline{\mathcal{W}}_1$ whose minimal signal is as large as possible. Given a consistent estimator $\hat{\Theta}^{(1)}$ and sufficiently strong signal strength, the solution of (3.5) identifies a minimal substructure of $G(\Theta^*)$ belonging to \mathcal{G}_1 . This strategy is motivated by the definition of the alternative parameter set $\mathcal{S}_1(\theta, s)$ (3.2). We remark that solving program (3.5) could be computationally challenging for some combinatorial tests. However, for all examples considered in this paper, simple and efficient polynomial time algorithms are available. We refer to the graph (\overline{V}, \hat{E}) as the *minimal structure witnessing the alternative*. Although the minimal structure witness is defined in full generality, for the ease of presentation we justify its validity on a case by case basis.

In the second step, NAWT attempts to certify the witness structure using the estimate $\hat{\Theta}^{(2)}$. Formally, NAWT aims to test the hypothesis

$$\mathbf{H}_0 : \exists e \in \hat{E} : \Theta_e^* = 0 \text{ vs } \mathbf{H}_1 : \forall e \in \hat{E} : \Theta_e^* \neq 0. \quad (3.6)$$

A rejection of the null hypothesis in (3.6) certifies the presence of the alternative witness structure. If the test fails to reject, NAWT cannot reject the null structure hypothesis. In Section 3.1 we give a detailed description on how NAWT works.

3.1 Minimal Structure Certification

In this section we detail an algorithm for testing (3.6). In fact, we present a general test for the following multiple testing problem

$$\mathbf{H}_0 : \exists e \in E \text{ s.t. } \Theta_e^* = 0 \text{ vs } \mathbf{H}_1 : \forall e \in E, \Theta_e^* \neq 0,$$

using the data $\mathcal{X} = \{\mathbf{X}_i\}_{i=1}^n$, $\mathbf{X}_i \sim N(0, (\Theta^*)^{-1})$, where E is a pre-given edge set. Following Neykov et al. (2015), for any $j, k \in [d]$ we define the bias corrected estimate:

$$\tilde{\Theta}_{jk} := \hat{\Theta}_{jk} - \frac{\hat{\Theta}_{*j}^T (\hat{\Sigma} \hat{\Theta}_{*k} - \mathbf{e}_k)}{\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j}},$$

where \mathbf{e}_k is a canonical unit vector with 1 at its k^{th} entry. Under mild regularity conditions, Neykov et al. (2015) show that $\hat{\Theta}_{jk}$ admits the following Bahadur representation:

$$\sqrt{n} \tilde{\Theta}_{jk} = n^{-1/2} \sum_{i=1}^n (\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{jk}^*) + o_p(1).$$

This motivates a multiplier bootstrap scheme for approximating the distribution of the statistic $\sqrt{n} \tilde{\Theta}_{jk}$ under the null hypothesis:

$$\widehat{W}_{jk} = n^{-1/2} \sum_{i=1}^n (\hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*k} - \hat{\Theta}_{jk}) \zeta_i,$$

where $\zeta_i \sim N(0, 1)$, $i \in [n]$ are independent and identically distributed. To approximate the null distribution of the statistic $\max_{(j,k) \in S} \sqrt{n} |\tilde{\Theta}_{jk}|$ over a subset $S \subseteq E$, let $c_{1-\alpha, S}$ denote the $(1 - \alpha)$ -quantile of the statistic $\max_{(j,k) \in S} |\widehat{W}_{jk}|$ (conditioned on the dataset \mathcal{X}). Formally, we let

$$c_{1-\alpha, S} := \inf_{t \in \mathbb{R}} \left\{ t : \mathbb{P} \left(\max_{(j,k) \in S} |\widehat{W}_{jk}| \leq t \mid \mathcal{X} \right) \geq 1 - \alpha \right\}. \quad (3.7)$$

As defined $c_{1-\alpha, S}$ is a population quantity (conditioning on \mathcal{X}). In practice an arbitrarily accurate estimate $\widehat{c}_{1-\alpha, S}$ of $c_{1-\alpha, S}$ can be obtained via Monte Carlo simulations. Below we describe a multiple edge testing procedure returning a subset $\widehat{E}^{\text{nc}} \subseteq E$ of rejected edges (hypotheses). Our procedure is based on the step-down construction of [Chernozhukov et al. \(2013\)](#), which is inspired by the multiple testing method of [Romano and Wolf \(2005\)](#).

Algorithm 1 Multiple Edge Testing

```

Initialize  $\widehat{E}^{\text{n}} \leftarrow E$ ;
repeat
  Reject  $R \leftarrow \{e \in \widehat{E}^{\text{n}} : \sqrt{n} |\tilde{\Theta}_e| \geq c_{1-\alpha, \widehat{E}^{\text{n}}}\}$ 
  Update  $\widehat{E}^{\text{n}} \leftarrow \widehat{E}^{\text{n}} \setminus R$ 
until  $R = \emptyset$  or  $\widehat{E}^{\text{n}} = \emptyset$ 
return  $\widehat{E}^{\text{nc}} \leftarrow E \setminus \widehat{E}^{\text{n}}$ 

```

Decompose the edge set $E = E^{\text{n}} \cup E^{\text{nc}}$, where $E^{\text{n}} \cap E^{\text{nc}} = \emptyset$, E^{n} is the subset of true null edges and E^{nc} is the set of non-null edges. Define the parameter space

$$\mathcal{M}_{E^{\text{n}}, E^{\text{nc}}}(s, \kappa) := \{\Theta \in \mathcal{M}(s) : \max_{e \in E^{\text{n}}} |\Theta_e| = 0, \min_{e \in E^{\text{nc}}} |\Theta_e| \geq \kappa \sqrt{\log d/n}\}. \quad (3.8)$$

For a fixed edge set E , we say that an edge set \widehat{E}^{r} has strong control of the family-wise error rate (FWER) if

$$\limsup_{n \rightarrow \infty} \sup_{E^{\text{n}} \subseteq E} \sup_{\Theta^* \in \mathcal{M}_{E^{\text{n}}, E^{\text{nc}}}(s, \kappa)} \mathbb{P}(E^{\text{n}} \cap \widehat{E}^{\text{r}} \neq \emptyset) \leq \alpha, \quad (3.9)$$

for some pre-specified size $\alpha > 0$. Our next result shows that Algorithm 1 returns an edge set \widehat{E}^{nc} with strong control of the FWER.

Proposition 3.1 (Strong FWER Test). Let $\Theta^* \in \mathcal{M}_{E^{\text{n}}, E^{\text{nc}}}(s, \kappa)$ and furthermore:

$$s \log(nd) \sqrt{\log d \log(nd)} / \sqrt{n} = o(1), \quad (\log(dn))^6 / n = o(1). \quad (3.10)$$

Then for any fixed edge set E , the output \widehat{E}^{nc} of Algorithm 1 satisfies (3.9). In addition, if the constant κ in (3.8) satisfies $\kappa \geq C' C^4$ for a fixed absolute constant $C' > 0$, we have:

$$\liminf_{n \rightarrow \infty} \inf_{E^{\text{n}} \subseteq E} \inf_{\Theta^* \in \mathcal{M}_{E^{\text{n}}, E^{\text{nc}}}(s, \kappa)} \mathbb{P}(E^{\text{nc}} = \widehat{E}^{\text{nc}}) = 1.$$

Of note, when κ is sufficiently large, Algorithm 1 achieves exact control of the FWER, i.e., we have equality in (3.9). This happens since all edges in E^{nc} will be rejected with overwhelming probability, while the bootstrap comparison is asymptotically exact for the remaining edges E^{n} . As a consequence of this result, if the null hypothesis set $\mathcal{S}_0(s)$ considered in (3.1) exhibits signal strength as in definition (2.2), the NAWT tests are exact.

Definition 3.1. For an edge set E let \widehat{E}^{nc} be the output of Algorithm 1 on data \mathcal{X} with level $1 - \alpha$. Define the following test function:

$$\psi_{\alpha,E}^B(\mathcal{X}) := \mathbb{1}(E = \widehat{E}^{\text{nc}}).$$

$\psi_{\alpha,E}^B(\mathcal{X})$ tests whether the set E is comprised only of non-null edges.

3.2 Examples

In this section we describe practical algorithms based on the NAWT for testing problems outlined in Section 2.1. Our tests can distinguish the null from the alternative hypotheses when the minimum signal strength is sufficiently large. As we shall see, the magnitude of the required signal strength is precisely of order $\sqrt{\log d/n}$ and therefore in view of Section 2 these tests are minimax optimal. Recall that we observe n i.i.d. samples $\{\mathbf{X}_i\}_{i=1}^n$ from $\mathbf{X}_i \sim N_d(0, (\boldsymbol{\Theta}^*)^{-1})$. We split the data into $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}$, $\mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$ and obtain estimates $\widehat{\boldsymbol{\Theta}}^{(1)}, \widehat{\boldsymbol{\Theta}}^{(2)}$ on \mathcal{D}_1 and \mathcal{D}_2 correspondingly, for which (3.3) and (3.4) hold.

3.2.1 Connectivity Testing

This section proposes a new procedure for honestly testing whether $G(\boldsymbol{\Theta}^*)$ is a connected graph. Accordingly, the sub-decomposition is $\mathcal{G}_0 := \{G \in \mathcal{G} : G \text{ disconnected}\}$ and $\mathcal{G}_1 := \{G \in \mathcal{G} : G \text{ connected}\}$. It is simple to check that definitions (3.1) and (3.2) reduce to:

$$\mathcal{S}_0(s) = \{\boldsymbol{\Theta} \in \mathcal{M}(s) : G(\boldsymbol{\Theta}) \in \mathcal{G}_0\}, \quad (3.11)$$

$$\mathcal{S}_1(\theta, s) = \{\mathcal{M}(s) : G(\boldsymbol{\Theta}) \in \mathcal{G}_1, \exists E \subset E(G(\boldsymbol{\Theta})), (\overline{V}, E) \text{ is a tree, } \min_{e \in E} |\Theta_e| \geq \theta\}, \quad (3.12)$$

Finding the minimal structure witness (3.5) reduces to finding a maximum spanning tree (MST) \widehat{T} on the full graph with edge weights $|\widehat{\Theta}_e^{(1)}|$. The complexity of finding a MST is $O(d^2 \log d)$, where d is the number of vertices. We summarize the procedure below:

Algorithm 2 Connectivity Test

Input: $\mathcal{D} = \{\mathbf{X}_i\}_{i=1}^n$, level $0 < \alpha < 1$.
Split the data $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}$, $\mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$
Using \mathcal{G}_1 obtain estimate $\widehat{\boldsymbol{\Theta}}^{(1)}$ satisfying (3.3)
Find MST \widehat{T} on the full graph with weights $|\widehat{\Theta}_e^{(1)}|$.
Output: $\psi_{\alpha, \widehat{T}}^B(\mathcal{D}_2)$.

The results on connectivity testing are summarized in the following

Corollary 3.2. Let $\theta = (2K + C'C^4)\sqrt{\log d/\lfloor n/2 \rfloor}$ for an absolute constant $C' > 0$ and assume that (3.10) holds. Then for any fixed level α , the test $\psi_{\alpha, \hat{T}}^B(\mathcal{D}_2)$ from the output of Algorithm 2 satisfies:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(s)} \mathbb{P}(\text{reject } \mathbf{H}_0) \leq \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\text{reject } \mathbf{H}_0) = 1.$$

3.2.2 Connected Components Testing

Connected component test is more general than the connectivity test. For an integer $\mathbf{m} \in [d-1]$ let $\mathbf{H}_0 : \# \text{ connected components} \geq \mathbf{m} + 1$ vs $\mathbf{H}_1 : \# \text{ connected components} \leq \mathbf{m}$. Testing connectivity is a special case when $\mathbf{m} = 1$.

Recall that a sub-graph F of G , i.e., $E(F) \subset E(G)$ and $V(F) = V(G) = \bar{V}$, is called a *spanning forest* of G , if F contains no cycles, adding any edge $e \in E(G) \setminus E(F)$ to $E(F)$ creates a cycle, and $|E(F)|$ is maximal. This definition extends naturally to graphs with positive weights on their edges. For a graph G let $\text{SF}(G)$ denote the class of spanning forests of G . Define the set of precision matrices $\mathcal{SF}(\theta) := \{\Theta : \exists F \in \text{SF}(G(\Theta)) \text{ such that } \min_{e \in F} |\Theta_e| \geq \theta\}$. For a fixed $\mathbf{m} \in [d-1]$ definitions (3.1) and (3.2) reduce to:

$$\begin{aligned} \mathcal{S}_0(s) &:= \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_{\mathbf{m}+1}\}, \\ \mathcal{S}_1(\theta, s) &:= \{\Theta \in \mathcal{SF}(\theta) \cap \mathcal{M}(s) : G(\Theta) \in \cup_{j \leq \mathbf{m}} \mathcal{G}_j\}, \end{aligned}$$

where $\mathcal{G}_j := \{G \in \mathcal{G} : G \text{ has exactly } j \text{ connected components}\}$ for all $j \in \mathbb{N}$.

Algorithm 3 Connected Components Test

Input: $\mathcal{D} = \{\mathbf{X}_i\}_{i=1}^n$, level $0 < \alpha < 1$.

Split the data into $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}$, $\mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$

Using \mathcal{D}_1 obtain an estimate $\hat{\Theta}^{(1)}$ satisfying (3.3)

Find an MSF \hat{F} with \mathbf{m} connected components with weights $|\hat{\Theta}_e^{(1)}|$, let $\hat{E} = E(\hat{F})$.

Output: $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$.

We have the following result regarding the asymptotic performance of Algorithm 3

Corollary 3.3. Let $\theta = (2K + C'C^4)\sqrt{\log d/\lfloor n/2 \rfloor}$ for an absolute constant $C' > 0$ and assume that (3.10) holds. Then for any fixed level α , the test $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$ from the output of Algorithm 3 satisfies:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(s)} \mathbb{P}(\text{reject } \mathbf{H}_0) \leq \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\text{reject } \mathbf{H}_0) = 1.$$

3.2.3 Cycle Detection

In this section we outline an algorithm testing whether the graph is acyclic. Recall the sub-decomposition $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ contains no cycles}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ contains a cycle}\}$. The parameter spaces (3.1) and (3.2) reduce to

$$\begin{aligned}\mathcal{S}_0(s) &:= \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_0\}, \\ \mathcal{S}_1(\theta, s) &:= \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_1, \exists E \subset E(G(\Theta)), E \text{ is a cycle, } \min_{e \in E} |\Theta_e| \geq \theta\}.\end{aligned}$$

The cycle presence test is compactly summarized below:

Algorithm 4 Cycle Presence Tests

Input: $\mathcal{D} = \{\mathbf{X}_i\}_{i=1}^n$, level $0 < \alpha < 1$.

Split the data into $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}$, $\mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$

Using \mathcal{D}_1 obtain estimate $\hat{\Theta}^{(1)}$ satisfying (3.3); Sort edges according to weights $|\hat{\Theta}_e^{(1)}|$

Iteratively add edges e from high to low $|\hat{\Theta}_e^{(1)}|$ values until a cycle is present.²

Let \hat{E} be set of edges forming the cycle.

Output: $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$.

The following result justifies the validity of Algorithm 4.

Corollary 3.4. Let $\theta = (2K + C'C^4)\sqrt{\log d / \lfloor n/2 \rfloor}$ for an absolute constant $C' > 0$ and assume that (3.10) holds. Then for any fixed level α , the test $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$ from the output of Algorithm 4 satisfies:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(s)} \mathbb{P}(\text{reject } \mathbf{H}_0) \leq \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\text{reject } \mathbf{H}_0) = 1.$$

3.2.4 Triangle-Free Graph Testing

In this section we discuss an algorithm for testing whether the graph is triangle-free. The corresponding sub-decomposition is $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ is triangle-free}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : \exists \text{ 3-clique subgraph of } G\}$. The parameter sets of precision matrices specified by (3.1) and (3.2) reduce to

$$\begin{aligned}\mathcal{S}_0(s) &:= \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_0\}, \\ \mathcal{S}_1(\theta, s) &:= \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_1, \exists E \subset E(G(\Theta)), E \text{ is a triangle, } \min_{e \in E} |\Theta_e| \geq \theta\}.\end{aligned}$$

We summarize the cycle test below

²This action requires no more than d iterations.

Algorithm 5 Triangle-Free Test

Input: $\mathcal{D} = \{\mathbf{X}_i\}_{i=1}^n$, level $0 < \alpha < 1$.

Split the data into $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}$, $\mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$

Using \mathcal{D}_1 obtain estimate $\widehat{\Theta}^{(1)}$ satisfying (3.3); Sort edges according to weights $|\widehat{\Theta}_e^{(1)}|$

Iteratively add edges e from high to low $|\widehat{\Theta}_e^{(1)}|$ values until a triangle emerges.

Let \widehat{E} be the set of edges forming the triangle.

Output: $\psi_{\alpha, \widehat{E}}^B(\mathcal{D}_2)$

The following proposition characterizes the performance of $\psi_{\alpha, \widehat{E}}^B(\mathcal{D}_2)$.

Corollary 3.5 (Honest Triangle-Free Test). Let $\theta = (2K + C'C^4)\sqrt{\log d/\lfloor n/2 \rfloor}$ for an absolute constant $C' > 0$ and assume that (3.10) holds. Then for any fixed level α , the test $\psi_{\alpha, \widehat{E}}^B(\mathcal{D}_2)$ from the output of Algorithm 6 satisfies:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(s)} \mathbb{P}(\text{reject } \mathbf{H}_0) \leq \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\text{reject } \mathbf{H}_0) = 1.$$

3.2.5 Self-Avoiding Path Length Testing

In this section we summarize an algorithm for testing whether the longest SAP in the graph is of length $\leq \mathbf{m}$ vs the longest SAP in the graph has length $\geq \mathbf{m} + 1$. The sub-decomposition is $\mathcal{G}_0 = \{G \in \mathcal{G} : \text{All SAPs in } G \text{ have length } \leq \mathbf{m}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : \exists \text{ a SAP of length } \mathbf{m} + 1\}$. The parameter sets of precision matrices specified by (3.1) and (3.2) reduce to:

$$\mathcal{S}_0(s) := \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_0\},$$

$$\mathcal{S}_1(\theta, s) := \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_1, \exists E \subset E(G(\Theta)), E \text{ is SAP of length } \mathbf{m} + 1, \min_{e \in E} |\Theta_e| \geq \theta\}.$$

The SAP length test is summarized below:

Algorithm 6 SAP Length Test

Input: $\mathcal{D} = \{\mathbf{X}_i\}_{i=1}^n$, level $0 < \alpha < 1$.

Split the data into $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}$, $\mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$

Using \mathcal{D}_1 to obtain estimate $\widehat{\Theta}^{(1)}$ satisfying (3.3); Sort edges according to weights $|\widehat{\Theta}_e^{(1)}|$

Iteratively add edges e from high to low $|\widehat{\Theta}_e^{(1)}|$ values until a SAP of length $\mathbf{m} + 1$ emerges.

Let \widehat{E} be the set of edges forming the SAP of length $\mathbf{m} + 1$.

Output: $\psi_{\alpha, \widehat{E}}^B(\mathcal{D}_2)$

Below we characterize the performance of the output of Algorithm 6 — $\psi_{\alpha, \widehat{E}}^B(\mathcal{D}_2)$.

Corollary 3.6 (Honest SAP Test). Let $\theta = (2K + C'C^4)\sqrt{\log d/\lfloor n/2 \rfloor}$ for an absolute constant $C' > 0$ and assume that (3.10) holds. Then for any fixed level α , the test $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$ from the output of Algorithm 6 satisfies:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(s)} \mathbb{P}(\text{reject } \mathbf{H}_0) \leq \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\text{reject } \mathbf{H}_0) = 1.$$

3.3 Extensions to More General Graphical Models

The NAWT procedure can be applied far beyond the class of Gaussian graphical models. Our first extension is the transelliptical graphical models proposed by Liu et al. (2012). We start with some definitions.

Definition 3.2 (Elliptical distribution Fang et al. (1990)). Let $\boldsymbol{\mu}^* \in \mathbb{R}^d$ and $\boldsymbol{\Sigma}^* \in \mathbb{R}^{d \times d}$. A d -dimensional random vector \mathbf{X} has an elliptical distribution, denoted with $\mathbf{X} \sim EC_d(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*, \xi)$, if $\mathbf{X} \stackrel{D}{=} \boldsymbol{\mu}^* + \xi \mathbf{A} \mathbf{U}$, where $\mathbf{U} \in \mathbb{R}^q$ is a random vector uniformly distributed on the unit sphere, $\xi \geq 0$ is a scalar random variable independent of \mathbf{U} , and $\mathbf{A} \in \mathbb{R}^{d \times q}$ is a deterministic matrix such that $\mathbf{A} \mathbf{A}^T = \boldsymbol{\Sigma}^*$.

Definition 3.3 (Transelliptical Distribution Liu et al. (2012)). A continuous random vector $\mathbf{X} = (X_1, \dots, X_d)^T$ has a transelliptical distribution, if there exist monotone univariate functions f_1, \dots, f_d and a non-negative random variable ξ , with $\mathbb{P}(\xi = 0) = 0$, such that:

$$(f_1(X_1), \dots, f_d(X_d))^T \sim EC_d(0, \boldsymbol{\Sigma}^*, \xi),$$

where $\boldsymbol{\Sigma}^*$ is symmetric with $\text{diag}(\boldsymbol{\Sigma}^*) = 1$ and $\boldsymbol{\Sigma}^* > 0$.

A graphical model structure over the family of transelliptical distributions is defined through the notion of the “latent generalized concentration matrix” — $\boldsymbol{\Theta}^* = (\boldsymbol{\Sigma}^*)^{-1}$. An edge is present between the two variables X_j, X_k if and only if $\Theta_{jk}^* \neq 0$. To construct an estimate of $\boldsymbol{\Theta}^*$, one first estimates $\boldsymbol{\Sigma}^*$ by a ranked based estimator $\hat{\boldsymbol{\Sigma}}$. Next one obtains a consistent estimator of $\hat{\boldsymbol{\Theta}}$ satisfying (3.3) and (3.4) by applying CLIME on $\hat{\boldsymbol{\Sigma}}$ (see Liu et al. (2012) for details). A decorrelation method can then be applied to substitute the estimates $\hat{\Theta}_{jk}$ with $\tilde{\Theta}_{jk}$, where $\tilde{\Theta}_{jk}$ are guaranteed to have asymptotic normal distributions (see Neykov et al. (2015) e.g). The NAWT immediately applies to this setup.

Our second extension is the semiparametric exponential family graphical model proposed by Yang et al. (2014). Below by indexing a vector $\mathbf{v} = (v_1, \dots, v_d)^T$ with $\mathbf{v}_{\setminus j}$ we mean omitting the j^{th} element of that vector.

Definition 3.4 (Semiparametric Exponential Family Distribution). A d -dimensional random vector $\mathbf{X} = (X_1, \dots, X_d) \in \mathbb{R}^d$ follows a semiparametric exponential family distribution if for any $j \in [d]$, the conditional density of $X_j | \mathbf{X}_{\setminus j}$ satisfies

$$p(x_j | \mathbf{x}_{\setminus j}) = \exp \left[x_j (\boldsymbol{\beta}_j^{*T} \mathbf{x}_{\setminus j}) + f_j(x_j) - b_j(\boldsymbol{\beta}_j, f_j) \right], \quad (3.13)$$

where $f_j(\cdot)$ is a base measure and $b_j(\cdot, \cdot)$ is the log-partition function.

The above model is *semiparametric* since the base measure functions f_j are unknown, and it forms a rich family of models including Gaussian, Poisson and Ising models. It is easy to see that $\beta_{jk} = 0$ if and only if X_j and X_k are conditionally independent. Thus the semiparametric exponential family naturally defines a graphical model. A pairwise rank based procedure has been proposed by Yang et al. (2014) to obtain an estimator $\hat{\beta}_j$ of β_j for all $j \in [d]$. Yang et al. (2014) also provide finite sample bounds on $\|\hat{\beta}_j - \beta_j\|_2$, $\|\hat{\beta}_j - \beta_j\|_1$ with the optimal rates of convergence. They further propose a decorrelated estimator $\tilde{\beta}_{jk}$ satisfying $\sqrt{n}(\tilde{\beta}_{jk} - \beta_{jk}) \rightsquigarrow N(0, \sigma_{jk}^2)$ for some limiting variance σ_{jk}^2 . Therefore, we can apply the NAWT procedure to obtain valid combinatorial structure tests.

4 Multi-Edge NAS

The results of Section 2 have two major limitations. Firstly, the null base G_0 is assumed to be of bounded degree. Secondly, our results cover only tests for which there exists a single-edge NAS. In this section we relax both of these conditions. The following examples illustrate relevant testing problems which do not fall into the framework of Section 2.

Maximum Degree Testing. Consider testing whether the maximum degree of the graph d_{\max} satisfies $d_{\max} \leq s_0$ vs $d_{\max} \geq s_1$, where $s_0 < s_1 \leq s$ are integers which are allowed to scale with n . In this case it is impossible to simultaneously construct a null base graph G_0 of bounded degree and a single-edge NAS set \mathbf{C} .

Clique Detection. Consider testing whether the graph is empty versus the alternative that the graph is an s -clique. In this test the null base is an empty graph, while the NAS set is an s -clique, and therefore cannot be a single-edge NAS.

To handle multiple edge NAS, we first extend Definitions 2.1 and 2.2 to allow for the above examples.

Definition 4.1 (Null-Alternative Separator Set). Let $G_0 = (\bar{V}, E_0) \in \mathcal{G}_0$ be a fixed graph under the null with adjacency matrix \mathbf{A}_0 . We call a collection of edge sets \mathbf{C} a (multi-edge) NAS with null base G_0 if for all edge sets $S \in \mathbf{C}$ we have $S \cap E_0 = \emptyset$ and $(\bar{V}, E_0 \cup S) \in \mathcal{G}_1$. For any edge set $S \in \mathbf{C}$, we denote the adjacency matrix of the graph (\bar{V}, S) with \mathbf{A}_S .

Definition 4.2 (Edge Set Geodesic Predistance). For two edge sets S and S' and a given graph G let:

$$d_G(S, S') = \min_{e \in S, e' \in S'} d_G(e, e').$$

We provide two generic strategies to obtain combinatorial inference lower bounds on the signal strength. The first strategy, described in Section 4.1, assumes that all $S \in \mathbf{C}$ satisfy

$|S| \leq U$ for some fixed constant U . The second strategy, presented in Section 4.2, does not require bounded cardinality of the edge sets S , but requires the null bases and NAS sets have some special combinatorial properties (More details will be provided later).

4.1 Bounded Edge Sets

Below we consider an extension of Theorem 2.1 for multi-edge NAS, where the number of edges in each set $S \in \mathcal{C}$ satisfy $|S| \leq U$ for some fixed integer $U \in \mathbb{N}$. In contrast to Section 2, here the graph G_0 is allowed to have unbounded degree.

Theorem 4.1. Let $G_0 \in \mathcal{G}_0$ be a graph under the null, and let \mathcal{C} be a multi-edge NAS with null base G_0 . Suppose that $n \geq (\log |\mathcal{C}|)/c_0$ for some sufficiently small $\kappa, c_0 > 0$, and that:

$$\theta \leq \kappa \sqrt{\frac{M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|)}{nU}} \wedge \frac{\kappa}{U(\|\mathbf{A}_0\|_2 + 2U)} \wedge \frac{1 - C^{-1}}{4(\|\mathbf{A}_0\|_1 + 2U)}, \quad (4.1)$$

Assuming $M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|) \rightarrow \infty$ as $n \rightarrow \infty$ we have:

$$\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1.$$

Theorem 4.1 is an extension of Theorem 2.1. Specifically, Theorem 2.1 corresponds the setting where $U = 1$, and $\|\mathbf{A}_0\|_2 \leq \|\mathbf{A}_0\|_1 \leq D$ (recall that D is an upper bound of the graph degree). Even though by assumption $U = \max_{S \in \mathcal{C}} |S|$ is bounded, we explicitly keep the dependency on U in (4.1) to reflect how the bound changes if U is allowed to scale. The first term on the right hand side of (4.1) is the structural packing entropy, while the remaining two terms ensure the parameter θ is small enough to construct a valid packing set (More details are provided in the proof).

We illustrate the usefulness of Theorem 4.1 by an example similar to the ones in Section 2.1. Consider testing whether the maximum degree of the graph $G(\Theta^*)$ is at most s_0 vs it is at least s_1 , where $s_0 < s_1 \leq s$ can increase with n but the null-alternative gap $s_1 - s_0$ remains bounded. Therefore we cannot apply Theorem 2.1 but should use Theorem 4.1 instead. Define the sub-decomposition $\mathcal{G}_0 = \{G : d_{\max}(G) \leq s_0\}$ and $\mathcal{G}_1 = \{G : d_{\max}(G) \geq s_1\}$ respectively.

Example 4.1 (Maximum Degree Test with Bounded Null-Alternative Gap). Let $\mathcal{S}_0(\theta, s)$ and $\mathcal{S}_1(\theta, s)$ be defined in (2.2) and (2.3). Assume that $s \log d/n = o(1)$, $s\sqrt{\log d/n} = O(1)$ and $s = O(d^\gamma)$ for some $\gamma < 1$. Then if κ is small enough and $\theta < \kappa\sqrt{\log d/n}$, we have

$$\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1.$$

In order to show the above result, we start by building a graph G_0 by constructing $\lfloor \frac{d}{s_1+1} \rfloor$ non-intersecting s_0 -star graphs first. Define the star graph centers $C_j = (s_1 + 1)j + 1$ for

$j = 0, \dots, \lfloor \frac{d}{s_1+1} \rfloor - 1$. Next, define G_0 by

$$G_0 := \left(\overline{V}, \bigcup_{j=0}^{\lfloor \frac{d}{s_1+1} \rfloor - 1} \bigcup_{k=1}^{s_0} \{(C_j, C_j + k)\} \right).$$

To define the NAS \mathbf{C} , simply connect each of the vertices C_j to the remaining $s_1 - s_0$ vertices in the block as follows:

$$\mathbf{C} := \left\{ \bigcup_{k=s_0+1}^{s_1} \{(C_j, C_j + k)\} \mid j = 0, \dots, \lfloor \frac{d}{s_1+1} \rfloor - 1 \right\}.$$




Figure 7: Test for the maximum degree $\mathcal{G}_0 := \{G : d_{\max}(G) \leq s_0\}$ vs $\mathcal{G}_1 = \{G : d_{\max}(G) \geq s_1\}$ with $s_0 = 3$, $s_1 = 5$ and $d = 18$. The solid edges represent $G_0 \in \mathcal{G}_0$ with maximum degree $s_0 = 3$. We construct the NAS $\mathbf{C} = \{(1, 5), (1, 6)\}, \{(7, 11), (7, 12)\}, \{(13, 17), (13, 18)\}$.

Since the predistance between any two different edge sets $S, S' \in \mathbf{C}$, is $d_{G_0}(S, S') = \infty$, the set \mathbf{C} itself is a $(\log |\mathbf{C}|)$ -packing set. The latter implies that $M(\mathbf{C}, d_{G_0}, \log |\mathbf{C}|) = \log |\mathbf{C}| \asymp \log(d/(s_1 + 1)) \asymp \log d$. In addition, one can easily check that $\|\mathbf{A}_0\|_2 = \sqrt{s_0}$ and $\|\mathbf{A}_0\|_1 = s_0$. By Theorem 4.1 we have that under the required scaling $\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1$, when $\theta < \kappa \sqrt{\log d/n}$ for a sufficiently small κ .

4.2 Scaling Edge Sets

Theorem 4.1 requires the cardinality of edge sets in the NAS \mathbf{C} to be bounded. In this section, we consider multi-edge NAS \mathbf{C} such that $S \in \mathbf{C}$ is allowed to increase with n . For this case, the previous notion of packing entropy based on geodesic predistance is no longer effective to characterize the lower bound. Instead, we introduce a new mechanism called buffer entropy to quantify the lower bound under scaling multi-edge NAS.

We first intuitively explain why the structural entropy in Theorem 4.1 may not be enough for handling NAS with scaling edge sets. Recall that Theorem 4.1 uses the structural entropy $M(\mathbf{C}, d_{G_0}, \log |\mathbf{C}|)$ to characterize the lower bound and the packing entropy is calculated




Figure 8: In all plots S , S' and G_0 are plotted with dashed, dotted and solid edges respectively (only (c) and (d) have solid edges). For both case (a) and case (b), we have $d_{G_0}(S, S') = 0$. However, the sets S and S' in case (a) should be “closer” to each other compared to those on case (b) as they have more overlapping vertices, but d_{G_0} fails to tell this difference. In case (c) and (d) for instance, we also have $d_{G_0}(S, S') = 1$ in both cases (c) and (d) and hence d_{G_0} cannot differentiate the two cases. However, in case (c), S has more vertices connected to S' than in case (d). Therefore, S and S' should also be “closer” in case (c) than case (d).

based on the edge set geodesic predistance d_{G_0} in Definition 4.2. Intuitively, for any $S, S' \in \mathbf{C}$, the smaller $d_{G_0}(S, S')$ is, the easier it is to perform a test. However, under the setting of NAS with scaling edge sets, different edge sets $S, S' \in \mathbf{C}$ may have multiple overlapping vertices and the notion of geodesic predistance is no longer precise enough to reflect the closeness between S and S' . For example, in Figure 8, for both case (a) and case (b), we have $d_{G_0}(S, S') = 0$. However, the sets S and S' in Figure 8 (a) should be “closer” to each other compared to those on Figure 8 (b) as they have more overlapping vertices, but d_{G_0} fails to tell this difference. Similar example holds in the case of non-overlapping edge sets. In Figure 8 (c) and (d) for instance, we also have $d_{G_0}(S, S') = 1$ in both cases (c) and (d) and hence d_{G_0} cannot differentiate the two cases. However, in Figure 8 (c), S has more vertices connected to S' than in Figure 8 (d). Therefore, S and S' should also be “closer” in Figure 8 (c) than Figure 8 (d).

From the discussion above, we see that the graph geodesic predistance fails to characterize the closeness between two multi-edge sets. Looking more closely in Fig 8, we can identify critical vertices measuring the “closeness” between the edge sets S and S' . Intuitively, the critical vertices should be the vertices which all paths connecting S and S' pass through. For instance in case (a) of Fig 8, we have 3 such critical vertices whereas in case (b) we have 2 such vertices. Similarly, in case (c) we have 6 critical vertices while in case (d) we have 4 critical vertices. Below we introduce a new concept called vertex buffer, which helps to measure the closeness between edge sets S and S' more precisely than the geodesic predistance.

Definition 4.3 (Vertex Buffer). Let $G = (V, E)$ be a graph and $V_1, V_2 \subseteq V$ be two vertex sets. A set $\mathcal{V} \subseteq V_1 \cup V_2$ is called a buffer between V_1 and V_2 if any path on G connecting some vertex in V_1 to some vertex in V_2 must contain at least one vertex $v_1 \in \mathcal{V} \cap V_1$ and one




Figure 9: Visualization of one of the buffers in $\mathcal{V}(V_1, V_2, G)$. Here V_1, V_2 are plotted with red and blue respectively and G is in solid edges. The vertices in the buffer are marked in the dashed squares. We can see that all paths passing from V_1 to V_2 must contain at least one of these vertices in the buffer.

vertex $v_2 \in \mathcal{V} \cap V_2$. We use $\mathcal{V}(V_1, V_2, G)$ to denote the collection of all buffers between V_1 and V_2 .

We visualize an example of buffer in Figure 9. Using the notion of vertex buffer, we can define the new closeness measure between S and S' . Recall that $V(S)$ denotes the vertex set of the edge set S . For a NAS \mathbf{C} with null base G_0 and any two edge sets $S, S' \in \mathbf{C}$ let $\mathcal{V}_{S,S'} \in \mathcal{V}(V(S), V(S'), G(\mathbf{A}_{S,S'}))$, where we denote

$$\mathbf{A}_{S,S'} := \mathbf{A}_0 + \mathbf{A}_S + \mathbf{A}_{S'}. \quad (4.2)$$

The buffer size $|\mathcal{V}_{S,S'}|$ is a measure of closeness of the edge sets S and S' . For instance, we have $|\mathcal{V}_{S,S'}| = 3$ in Fig 8 (a) while $|\mathcal{V}_{S,S'}| = 2$ in Fig 8 (b) and we have $|\mathcal{V}_{S,S'}| = 6$ in Fig 8 (c) while $|\mathcal{V}_{S,S'}| = 4$ in Fig 8 (d). This implies that the larger $|\mathcal{V}_{S,S'}|$ is, the closer S and S' will be. The vertex buffer $\mathcal{V}_{S,S'}$ is useful for the lower bound since it characterizes the critical vertices between S and S' . The paths passing through at least one edge in both S and S' must pass through at least one vertex in $\mathcal{V}_{S,S'}$.

In contrast to the bounded edge sets case, when the edge sets in \mathbf{C} are allowed to scale in size, it is not effective to build packing sets based on the predistance, since this strategy limits the number of edge sets we can build. One way to increase the cardinality of \mathbf{C} is to consider a larger number of potentially overlapping structures, and use the buffer size as a more precise closeness measure between these structures. Below we introduce a new entropy concept which quantifies this intuition.

Definition 4.4 (Buffer Entropy). Let \mathbf{C} be a multi-edge NAS set with a base graph G_0 and $\mathbf{V} := \{\mathcal{V}_{S,S'}\}_{S,S' \in \mathbf{C}}$ be a collection of buffers. The buffer entropy is defined as:

$$M_B(\mathbf{C}, \mathbf{V}) := \log \left(\left[\max_{S \in \mathbf{C}} \mathbb{E}_{S'} |\mathcal{V}_{S,S'}| \right]^{-1} \right), \quad (4.3)$$

where the expectation $\mathbb{E}_{S'}$ is taken with respect to uniformly sampling S' from \mathbf{C} .

Under some cases, the buffer entropy $M_B(\mathbf{C}, \mathbf{V})$ is asymptotically equivalent to the packing entropy $M(\mathbf{C}, d_{G_0}, \log |\mathbf{C}|)$. For instance, in Example 4.1, we can see from Figure 7 that $\mathcal{V}_{S,S'} = \emptyset$ if $S \neq S'$ and $\mathcal{V}_{S,S'} = V(S)$ if $S = S'$. Therefore, for all $S \in \mathbf{C}$, $\mathbb{E}_{S'}[|\mathcal{V}_{S,S'}|] = (s_1 - s_0)/|\mathbf{C}|$ and the buffer entropy $M_B(\mathbf{C}, \mathbf{V}) = \log(|\mathbf{C}|/(s_1 - s_0)) \asymp M(\mathbf{C}, d_{G_0}, \log |\mathbf{C}|)$ since $s_1 - s_0 \leq U$ is a bounded constant. However, for edge sets of scaling cardinalities, the buffer entropy may have a better characterization of the complexity of NAS set than the packing entropy as we argued by the examples in Figure 8.

We want the buffer entropy to be as large as possible to achieve sharp lower bounds. Here we provide some guidelines on how to find the buffer sets. In general, we aim to find the buffers $\mathcal{V}_{S,S'}$ of size as small as possible for any pair $S, S' \in \mathbf{C}$. It is easy to see that the set $\{V(E_0 \cup S) \cap V(S')\} \cup \{V(E_0 \cup S') \cap V(S)\}$ is always a valid vertex buffer, and we believe it is a good choice for $\mathcal{V}_{S,S'}$ in general (see Fig 9 and Examples 4.2, 4.3, e.g.)

Since larger $|\mathcal{V}_{S,S'}|$ implies that S and S' are closer, we aim to build NAS in which the cardinality $|\mathcal{V}_{S,S'}|$ is not too large. We have the trivial bound

$$|\mathcal{V}_{S,S'}| \leq \sum_{v \in V(S)} \mathbf{1}(v \in \mathcal{V}_{S,S'}) + \sum_{v \in V(S')} \mathbf{1}(v \in \mathcal{V}_{S,S'}).$$

Intuitively, if the random variables $\{\mathbf{1}(v \in \mathcal{V}_{S,S'})\}_{v \in V(S)}$ are negatively correlated, $|\mathcal{V}_{S,S'}|$ will not be too large. We formalize this idea in the following definition.

Definition 4.5 (Incoherent NAS). The collection of edge sets \mathbf{C} is called an incoherent NAS under a collection of buffers $\mathbf{V} := \{\mathcal{V}_{S,S'}\}_{S,S' \in \mathbf{C}}$ if for any fixed $S \in \mathbf{C}$, the random variables $\{\mathbf{1}(v \in \mathcal{V}_{S,S'})\}_{v \in V(S)}$ with respect to a uniformly sampled S' from \mathbf{C} are negatively associated. In other words, for any pair of disjoint sets $I, J \subseteq V(S)$ and any pair of coordinate-wise nondecreasing functions f, g we have:

$$\text{Cov}(f(\{\mathbf{1}(v \in \mathcal{V}_{S,S'})\}_{v \in I}), g(\{\mathbf{1}(v \in \mathcal{V}_{S,S'})\}_{v \in J})) \leq 0.$$

Negative association is satisfied by a variety of classical discrete distributions such as the multinomial and hypergeometric, and even more generally by the class of permutation distributions (Joag-Dev and Proschan, 1983; Dubhashi and Ranjan, 1996, e.g.). It is a standard assumption that has been exploited in other works (Addario-Berry et al., 2010, e.g.) for obtaining lower bounds. We will also show concrete constructions of incoherent NAS in Examples 4.2 and 4.3.

Besides packing entropy, the lower bound in Theorem 4.1 also involves the maximum degree $\|\mathbf{A}_0\|_1$ and the spectral norm $\|\mathbf{A}_0\|_2$. We define similar quantities for the scaling edge sets case. As the sizes of $S, S' \in \mathbf{C}$ are no longer ignorable, we need to consider the matrix

$\mathbf{A}_{S,S'}$ (4.2) instead. Denote the uniform maximum degree as $\Gamma := \max_{S,S' \in \mathcal{C}} \|\mathbf{A}_{S,S'}\|_1$ and uniform spectral norm as $\Lambda := \max_{S,S' \in \mathcal{C}} \|\mathbf{A}_{S,S'}\|_2$. We define

$$\mathcal{R} := \max_{S,S' \in \mathcal{C}} \frac{|S \cap S'|}{|\mathcal{V}_{S,S'}|}, \quad \mathcal{B} := \max_{S,S' \in \mathcal{C}} \left((\Gamma^2 |\mathcal{V}_{S,S'}|) \wedge \Lambda^4 \right).$$

\mathcal{R} is an edge-node ratio measuring how dense the edge set $S \cap S'$ is compared to the vertex buffers. For example, in Figure 10, the ratio \mathcal{R} is 1 in case (a) compared to 1/2 in case (b). The quantity \mathcal{B} is an auxiliary quantity assembling maximum degrees, spectral norms and buffer sizes and helps to obtain a compact lower bound formulation.




Figure 10: Illustration of an edge-node ratios. For case (a), we have $\mathcal{R} = 1$ and $\mathcal{R} = 1/2$ for case (b). This implies that edges in case (a) are denser than case (b).

Below we connect the structural features we defined above to the lower bound. Recall definitions (2.2) and (2.3) on $\mathcal{S}_0(\theta, s)$ and $\mathcal{S}_1(\theta, s)$. We have the following theorem.

Theorem 4.2. Let \mathcal{C} be an incoherent NAS under the collection of buffers $\mathbf{V} := \{\mathcal{V}_{S,S'}\}_{S,S' \in \mathcal{C}}$. Then if $M_B(\mathcal{C}, \mathbf{V}) \rightarrow \infty$ and

$$\theta \leq \sqrt{\frac{M_B(\mathcal{C}, \mathbf{V})}{4n\mathcal{R}}} \wedge \sqrt{\frac{\mathcal{R}}{\mathcal{B}}} \wedge \frac{1 - C^{-1}}{2\sqrt{2}\Gamma}, \quad (4.4)$$

the minimax risk satisfies $\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1$.

When the sample size n is sufficiently large, the buffer entropy term on the right hand side of (4.4) is the smallest term and drives the bound which bares similarity to Theorem 4.1. As we previously argued, in Example 4.1, the buffer entropy and the structural entropy are of the same order, and moreover it is easy to check that the quantities \mathcal{R} and \mathcal{B} satisfy $\mathcal{R} = \frac{|S|}{|S|+1}$ and $\mathcal{B} \leq (U+1)(\|\mathbf{A}_0\|_1 + 2U)^2$ which implies that when $\max_{S \in \mathcal{C}} |S| \leq U$ the terms in bound (4.4) match those from (4.1) up to constants.

To better illustrate the usage of Theorem 4.2 we consider two examples. First we focus on the problem of testing whether the maximum degree in the graph is $\leq s_0$ vs $\geq s_1$. In contrast to Example 4.1, here s_0 is allowed to be much smaller than s_1 and we cannot apply Theorem 4.1. When $s_0 = 0$, this problem is related to the problem of detecting a set of s_1 signals in




Figure 11: Test for maximum degree $\mathcal{G}_0 = \{G : d_{\max}(G) \leq s_0\}$ and $\mathcal{G}_1 = \{G : d_{\max}(G) \geq s_1\}$ with $s_0 = 3$ and $s_1 = 6$. We split the vertices into two parts $\{1, \dots, \lfloor \sqrt{d} \rfloor\}$ and $\{\lfloor \sqrt{d} \rfloor + 1, \dots, d\}$. We use the first part of vertices to construct s_0 -star graphs as G_0 (visualized with solid edges). To construct the NAS \mathcal{C} , we select any $s_1 - s_0$ vertices (e.g., vertices 13, 14, 16) from the second vertices part $\{\lfloor \sqrt{d} \rfloor + 1, \dots, d\}$ and connect them to any center of the s_0 -star graphs in G_0 (e.g., vertex 1). This gives us one of the edge sets $S \in \mathcal{C}$ (e.g., $S = \{(1, 13), (1, 14), (1, 16)\}$ depicted in dashed edges in the figure). We construct \mathcal{C} containing all such edge sets. For any $S, S' \in \mathcal{C}$, we choose the vertex buffer $\mathcal{V}_{S, S'} = V(S) \cap V(S')$ (e.g., for $S = \{(1, 13), (1, 14), (1, 16)\}$ and $S' = \{(5, 14), (5, 16), (5, 18)\}$ we have $\mathcal{V}_{S, S'} = \{14, 16\}$ and $|S \cap S'| = \emptyset$.)

the normal means model (Ingster, 1982; Baraud, 2002; Donoho and Jin, 2004; Addario-Berry et al., 2010; Verzelen and Villers, 2010, e.g.). However the two problems are distinct, since we are studying structural testing in the graphical model setting. Given $s_0 < s_1 \leq s$, we let the sub-decomposition be $\mathcal{G}_0 = \{G : d_{\max}(G) \leq s_0\}$ and $\mathcal{G}_1 = \{G : d_{\max}(G) \geq s_1\}$. We summarize our results in the following

Example 4.2 (Maximum Degree Test with Scaling Null-Alternative Gap). Let $\mathcal{S}_0(\theta, s)$ and $\mathcal{S}_1(\theta, s)$ be defined in (2.2) and (2.3). Assume that $s\sqrt{\log d/n} = O(1)$ and $s = O(d^\gamma)$ for some $\gamma < 1/2$. Then for a small enough absolute constant κ if $\theta < \kappa\sqrt{\log d/n}$ we have $\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1$.

Before we show how this example follows from Theorem 4.2 we would like to highlight the difference between Examples 4.1 and 4.2. First, note that Example 4.2 is more flexible compared to Example 4.1 since the number of edges is allowed to scale with n . However Example 4.2 is also more restrictive in that we require $s = O(d^\gamma)$ for some $\gamma < 1/2$, which is not required by Example 4.1.

We now first explain the construction of (G_0, \mathcal{C}) and then we formalize it. We start by splitting the vertices into two parts $\{1, \dots, \lfloor \sqrt{d} \rfloor\}$ and $\{\lfloor \sqrt{d} \rfloor + 1, \dots, d\}$. The graph G_0 is constructed based only on the first part of vertices, and consists of non-intersecting s_0 -star graphs. To build the NAS \mathcal{C} , we select any $s_1 - s_0$ vertices from the second set vertices $\{\lfloor \sqrt{d} \rfloor + 1, \dots, d\}$ and connect them to any center of the s_0 -star graphs in G_0 (e.g., vertex 1). The set \mathcal{C} contains all such edge sets (see Figure 11). Formally, to construct a graph G_0

we split the vertex set $\{1, \dots, \lfloor \sqrt{d} \rfloor\}$ into $\lfloor \frac{\sqrt{d}}{s_0+1} \rfloor$ non-intersecting s_0 -star graphs. We take centers $C_j = (s_0 + 1)j + 1$ for $j = 0, \dots, \lfloor \frac{\sqrt{d}}{s_0+1} \rfloor - 1$ and connect them to the remaining vertices as follows:

$$G_0 := \left(\overline{V}, \bigcup_{j=0}^{\lfloor \frac{\sqrt{d}}{s_0+1} \rfloor - 1} \bigcup_{k=1}^{s_0} \{(C_j, C_j + k)\} \right).$$

Since $s_0 < s_1 \leq s = O(d^\gamma)$, this creates at least $\lfloor \sqrt{d}/d^\gamma \rfloor \asymp d^{1/2-\gamma}$ graphs. Denote the set of centers of the s_0 -star graphs with

$$\mathcal{J} := \bigcup_{j=0}^{\lfloor \frac{\sqrt{d}}{s_0+1} \rfloor - 1} \{C_j\}.$$

To construct the NAS set \mathbf{C} from the remaining $d - \lfloor \sqrt{d} \rfloor$ isolated vertices we choose a set I of $s_1 - s_0$ vertices and connect them with any of the vertex C from the center set \mathcal{J} . Formally we let $\mathbf{C} := \{S_{C,I}\}_{C \in \mathcal{J}, I \in \mathcal{I}}$, where

$$S_{C,I} = \{(C, i)\}_{i \in I} \text{ and } \mathcal{I} = \left\{ I \in \{M \in 2^{[d]} : |M| = s_1 - s_0\} \mid \min_{i \in I} i > \lfloor \sqrt{d} \rfloor \right\}.$$

Consider the collection of vertex buffers $\mathcal{V}_{S,S'} = V(S) \cap V(S')$. A visualization of the key quantities $G_0, \mathbf{C}, \mathcal{V}_{S,S'}$ for two specific $S, S' \in \mathbf{C}$ is provided on Figure 11. Now we will argue that the set $\{\mathbf{1}(v \in \mathcal{V}_{S,S'})\}_{v \in V(S)}$ is negatively associated, as required by Assumption 4.5. Note that uniformly selecting a set $S' \in \mathbf{C}$ is equivalent to uniformly selecting a center $C \in \mathcal{J}$ and a vertex set $I \in \mathcal{I}$ and construct S' by connecting C to every vertex in I . By construction the selection of a center C is independent of the selection of the vertex set I . By a result in Section 3.1 (c) of Joag-Dev and Proschan (1983) the random variables $\{\mathbf{1}(v \in \mathcal{V}_{S,S'})\}_{v \in V(S)}$ are negatively associated. Next we calculate the remaining quantities from Theorem 4.2. By the triangle inequality $\|\mathbf{A}_{S,S'}\|_2 \leq \sqrt{s_1} + \sqrt{s_1 - s_0} \leq 2\sqrt{s_1}$ and $\|\mathbf{A}_{S,S'}\|_1 \leq 2s_1$ and therefore $\Lambda \leq 2\sqrt{s_1}$, $\Gamma \leq 2s_1$, and $\mathcal{B} \leq \Lambda^4 \leq 16s_1^2 \leq 16s^2$. To calculate upper and lower bounds of \mathcal{R} we only need to consider $|S \cap S'| \neq 0$ and in this case S and S' must share the same center. In this situation by definition we have $|S \cap S'| = |\mathcal{V}_{S,S'}| - 1$ and therefore $1/2 \leq \mathcal{R} \leq 1$.

Finally we calculate $\max_{S \in \mathbf{C}} \mathbb{E}_{S'} |\mathcal{V}_{S,S'}|$. Recall that $\mathcal{V}_{S,S'} = V(S) \cap V(S')$ and that a set S' can be constructed by first selecting a center C uniformly from the set \mathcal{J} and then selecting a set I uniformly from \mathcal{I} . Hence for all $S \in \mathbf{C}$

$$\mathbb{E}_{S'} |\mathcal{V}_{S,S'}| = \mathbb{E}[\mathbf{1}(C \in V(S)) |S|] + \mathbb{E}[|I \cap V(S)| \mid S] = \frac{1}{|\mathcal{J}|} + \frac{(s_1 - s_0)^2}{d - (s_0 + 1)|\mathcal{J}|}.$$

The last equality implies that $M_B(\mathbf{C}, \mathbf{V}) \asymp \log d$. An application of Theorem 4.2, and taking into account the required scaling completes the result.




Figure 12: Test for sparse clique with $s = 5$. We construct $G_0 = (\bar{V}, \emptyset)$ and visualize two intersecting $S, S' \in \mathcal{C}$, where S is the 5-clique $K_{\{1,3,5,7,9\}}$, S' is the 5-clique $K_{\{1,2,3,9,10\}}$ and $\mathcal{V}_{S,S'} = \{1, 3, 9\}$.

Our second example further illustrates the usage of Theorem 4.2 with a clique detection problem. Define the null and alternative parameter spaces:

$$\mathcal{S}_0 := \{\mathbf{I}_d\}, \quad \mathcal{S}_1(\theta, s) := \{\mathbf{I}_d + \theta(\mathbf{v}\mathbf{v}^T - \mathbf{I}_d) : \theta \in (0, 1), \forall j \in [d] \ v_j \in \{\pm 1, 0\}, \|\mathbf{v}\|_2^2 = s\}.$$

This setup is related to that in Berthet and Rigollet (2013); Johnstone and Lu (2009). Our case is different from previous works because we parametrize the precision matrix rather than the covariance matrix, and the parametrization is distinct. Under our parametrization, the graph in the alternative hypothesis consists of a single s -clique.

Example 4.3 (Sparse Clique Detection). For values of $\theta < \frac{1}{4\sqrt{2}s} \wedge \sqrt{\frac{\log(d/s^2)}{4ns}}$ we have:

$$\liminf_{n \rightarrow \infty} \gamma(\mathcal{S}_0, \mathcal{S}_1(\theta, s)) = 1.$$

Consider $G_0 = (\bar{V}, \emptyset)$, and let $\mathcal{C} = \{\text{all } s \text{ cliques with vertices in } \bar{V}\}$. Take two sets $S, S' \in \mathcal{C}$. Define the vertex buffer as $\mathcal{V}_{S,S'} = V(S) \cap V(S')$. Similarly to Example 4.2, conditioning on S the random variables $\{\mathbf{1}(v \in \mathcal{V}_{S,S'})\}_{v \in V(S)}$ are negatively associated (see Section 3.1 (c) of Joag-Dev and Proschan (1983)). By the definition of vertex buffer sets we have that $|\mathcal{V}_{S,S'}| \leq s$, $|S \cap S'| = \binom{\mathcal{V}_{S,S'}}{2} < |\mathcal{V}_{S,S'}|^2 \leq s|\mathcal{V}_{S,S'}|$ and hence $s - 1/2 \leq \mathcal{R} \leq s$. Furthermore $\Gamma \leq 2s$, and therefore $\Lambda \leq 2s$. Therefore $\mathcal{B} \leq 4s^3$. For all $v \in V(S)$ we have $\mathbb{E}_{S'}[\mathbf{1}(v \in \mathcal{V}_{S,S'})] = s/d$ which implies $\mathbb{E}_{S'}[|\mathcal{V}_{S,S'}|] = s^2/d$ for all $S \in \mathcal{C}$, and therefore $M_B(\mathcal{C}, \mathbf{V}) = \log(d/s^2)$. An application of Theorem 4.2 shows the statement.

4.3 Upper Bounds

In this section we develop algorithms matching the bounds in Sections 4.1 and 4.2.

4.3.1 Maximum Degree Test

We propose a test matching lower bound in Example 4.1. Recall that the sub-decomposition is $\mathcal{G}_0 = \{G : d_{\max}(G) \leq s_0\}$ and $\mathcal{G}_1 = \{G : d_{\max}(G) \geq s_1\}$. The parameter sets (3.1) and (3.2) reduce to $\mathcal{S}_0(s) = \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_0\}$ and $\mathcal{S}_1(\theta, s) = \{\Theta \in \mathcal{M}(s) : G(\Theta) \in \mathcal{G}_1, \exists \text{ vertex } k \in \bar{V}, \text{ and set } S \subset [d] \text{ such that } k \notin S, |S| = s, \min_{j \in S} |\Theta_{jk}| \geq \theta\}$. Recall the definitions of $\mathcal{D}_1, \mathcal{D}_2, \hat{\Theta}^{(1)}$ and $\hat{\Theta}^{(2)}$ from Section 3.2.

Algorithm 7 Maximum Degree Test

Input: $\mathcal{D} = \{\mathbf{X}_i\}_{i=1}^n$, level $0 < \alpha < 1$.

Split the data into $\mathcal{D}_1 = \{\mathbf{X}_i\}_{i=1}^{\lfloor n/2 \rfloor}, \mathcal{D}_2 = \{\mathbf{X}_i\}_{i=\lfloor n/2 \rfloor+1}^n$

Add edges e from high to low $|\hat{\Theta}_e^{(1)}|$ values until a vertex with $(s_0 + 1)$ neighbors occurs.

Let \hat{E} be the set of s edges connecting the vertex to its $(s_0 + 1)$ neighbours.

Output: $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$.

Following the formulation of Algorithm 7 we characterize the performance of $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$.

Corollary 4.3 (Honest Max Degree Test). Let $\theta = (2K + C'C^4)\sqrt{\log d / \lfloor n/2 \rfloor}$ for an absolute constant $C' > 0$ and assume that (3.10) holds. Then for any fixed level α , the test $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$ from the output of Algorithm 7 satisfies:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(s)} \mathbb{P}(\text{reject } \mathbf{H}_0) \leq \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\text{reject } \mathbf{H}_0) = 1.$$

Remark 4.1. The test of Algorithm 7 can also be used to test the problem described in Example 4.2, since the lower bound on the signal strength is of the same magnitude.

4.3.2 Clique Detection Test

Here we devise a test to match the lower bound of Example 4.3. Unlike previous tests, the clique detection test is not computationally feasible. Recall that $\hat{\Sigma} = n^{-1} \sum_{i=1}^n \mathbf{X}_i^{\otimes 2}$. Define

$$\hat{\lambda}_{\min} = \min_{|C|=s} \lambda_d(\hat{\Sigma}_{CC}),$$

where $\hat{\Sigma}_{CC}$ is a sub-matrix of $\hat{\Sigma}$ with both column and row indices in C , and $\lambda_d(\mathbf{A})$ is the smallest eigenvalue of the matrix \mathbf{A} . Consider the test $\psi = \mathbb{1}(\hat{\lambda}_{\min} < \nu)$, where ν is

$$\nu := \left(1 - (\sqrt{2} + 1)\sqrt{(s \log(ed/s) + \log(2\alpha^{-1}))/n}\right)^2,$$

for some small constant $\alpha \geq 0$. We have

Proposition 4.4. Given any $\alpha \in (0, 1)$, suppose $(s \log(ed/s) + \log(2\alpha^{-1}))/n = o(1)$. Then for values of $\theta \in (0, 1)$ satisfying

$$\theta > \kappa \sqrt{\frac{\log(ed/s)}{sn}},$$

for an absolute constant κ , we have

$$\limsup_{n \rightarrow \infty} \sup_{\Theta \in \mathcal{S}_0} \mathbb{P}_{\Theta}(\text{reject } \mathbf{H}_0) \leq \alpha \quad \liminf_{n \rightarrow \infty} \inf_{\Theta \in \mathcal{S}_1(\theta, s)} \mathbb{P}_{\Theta}(\text{reject } \mathbf{H}_0) = 1.$$

The proof of Proposition 4.4 is deferred to Appendix C.2. We remark that the bound of Proposition 4.4 matches the lower bound of Example 4.3 up to a scalar in the regime $s = O(d^\gamma)$ for $\gamma < 1/2$, and $s \log d/n = o(1)$.

5 Numerical Studies and Real Data Analysis

In this section we present numerical analysis of Algorithms 2 and 4 for the hypothesis tests on the connectivity and cycles on synthetic datasets. In addition, we implement the graph connectivity test to study the brain networks for the ADHD-200 dataset.

5.1 Connectivity Testing

We present numerical simulations assessing the performance of Algorithm 2 for testing connectivity. Recall $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ disconnected}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ connected}\}$, and consider testing

$$\mathbf{H}_0 : G(\Theta^*) \in \mathcal{G}_0 \text{ vs } \mathbf{H}_1 : G(\Theta^*) \in \mathcal{G}_1.$$

In order to construct the connected and disconnected graphs, we consider a chain graph of length k with adjacency matrix as follows

$$\mathbf{A}_{\text{Chain}(k)} = \begin{pmatrix} 0 & 1 & & 0 \\ 1 & 0 & 1 & \\ & 1 & \ddots & \ddots \\ & & \ddots & \ddots & 1 \\ 0 & & & 1 & 0 \end{pmatrix} \in \mathbb{R}^{k \times k}. \quad (5.1)$$

We construct the connected graph with adjacency matrix $\mathbf{A} = \mathbf{A}_{\text{Chain}(d)}$ and the disconnected graphs with adjacency matrices

$$\mathbf{A}_m = \text{diag}(\mathbf{A}_{\text{Chain}(m)}, \mathbf{A}_{\text{Chain}(d-m)}), \text{ for } m \in [d-1].$$

The corresponding graphs of \mathbf{A} and $\mathbf{A}_1, \dots, \mathbf{A}_{d-1}$ are visualized in Figure 13.




Figure 13: $G(\mathbf{A}_4)$ and $G(\mathbf{A})$ for $d = 7$.

To show that our test is uniformly valid over various disconnected graphs under the null hypothesis, we use different precision matrices for different repetitions. For each repetition, we randomly select $M \sim \text{Uniform}([d-1])$, choose the precision matrix $\Theta_M^*(\theta) = \mathbf{I}_d + \theta \mathbf{A}_M$ and generate i.i.d. samples $\mathbf{X}_i \sim N(\mathbf{0}, (\Theta_M^*(\theta))^{-1})$ for $i \in [n]$. Under the alternative hypothesis, we generate n i.i.d. samples $\mathbf{X}_1, \dots, \mathbf{X}_n$ from $N(\mathbf{0}, (\Theta^*(\theta))^{-1})$, where $\Theta^*(\theta) = \mathbf{I}_d + \theta \mathbf{A}$.

The sample size is chosen from $n = 400$ and 600 and the dimension varies in $d = 100, 125$ and 150 . The significance level of the tests is set to $\alpha = 0.05$. Within each test we generate 3,000 bootstrap samples to estimate the quantiles (3.7). We estimate the precision matrix by the CLIME estimator (Cai et al., 2011) as follows

$$\hat{\Theta}_\lambda = \text{argmin } \|\Theta\|_1, \quad \text{s.t. } \|\hat{\Sigma}\Theta - \mathbf{I}_d\|_{\max} \leq \lambda. \quad (5.2)$$

The tuning parameter is chosen by minimizing a K -fold cross validation risk

$$\text{CV}(\lambda) = \sum_{k=1}^K \|\hat{\Sigma}^{(k)} \hat{\Theta}_\lambda^{(-k)} - \mathbf{I}_d\|_{\max},$$

where $\hat{\Sigma}^{(k)}$ is the sample covariance matrix estimated on the k -th subset of data and $\hat{\Theta}_\lambda^{(-k)}$ is the CLIME estimator without using the k -th subset of data. We use the first scenario and implement a 5-fold cross validation. It selects $\lambda = 1.5\sqrt{\log d/n}$ and we keep using this value of λ throughout all remaining settings. We vary the signal strength $\theta \in [0.25, 0.45]$ and estimate the risk of our test ψ given by Algorithm 2:

$$\mathbb{P}_{G(\Theta^*) \in \mathcal{G}_0}(\psi = 1) + \mathbb{P}_{G(\Theta^*) \in \mathcal{G}_1}(\psi = 0),$$

by averaging the type-I and type-II errors under null and alternative settings through 200 repetitions. The results are visualized in Figure 14, where the Y-axis represents the risk and the X-axis plots the signal strength θ . As predicted by our results, with the increase in θ the risk curves become well controlled about the target value 0.05. Moreover, there is little difference between the three risk curves within each plot. The curves for $d = 125, 150$ have a slightly worse performance than the curves for smaller dimension $d = 100$. This effect is expected, as the signal strength required to separate the null and alternative depends on d only logarithmically. Furthermore, smaller signal strength is required to distinguish the null from the alternative when we increase n from 400 to 600. Table 1 reports the type-I errors of ψ . The type-I errors remain within the desired range of 0.05. As we increase the signal

strength θ , the values of the errors increase from 0 to the target level. This phenomenon is expected and is explained below Proposition 3.1.




Figure 14: Type-I and type-II errors control for connectivity test with $\theta \in [0.25, 0.45]$. The grey horizontal line is the significance level $\alpha = 0.05$.

Table 1: Sizes for connectivity testing $\theta \in [0.25, 0.45]$

n	θ		0.25	0.28	0.32	0.35	0.38	0.42	0.45
	d								
400	100		0.000	0.000	0.000	0.020	0.045	0.045	0.060
	125		0.000	0.000	0.000	0.010	0.045	0.050	0.060
	150		0.000	0.000	0.000	0.010	0.040	0.050	0.055
600	100		0.000	0.005	0.020	0.045	0.030	0.025	0.030
	125		0.000	0.000	0.025	0.050	0.040	0.045	0.040
	150		0.000	0.000	0.015	0.055	0.040	0.045	0.040

5.2 Cycle Presence Testing

In this section we present numerical analysis of Algorithm 4 for cycle presence testing. Recall the sub-decomposition $\mathcal{G}_0 = \{G \in \mathcal{G} : G \text{ contains no cycles}\}$ and $\mathcal{G}_1 = \{G \in \mathcal{G} : G \text{ contains a cycle}\}$. We aim to conduct the hypothesis

$$\mathbf{H}_0 : G(\Theta^*) \in \mathcal{G}_0 \text{ vs } \mathbf{H}_1 : G(\Theta^*) \in \mathcal{G}_1.$$

Similarly to the previous section, we also generate data under both the null and alternative hypotheses. We use the chain graph with adjacency matrix $\mathbf{A}_{\text{Chain}(d)}$ as the acyclic graph and construct the loopy graphs by adding edges to the chain graph. More specifically, we construct the graphs with adjacency matrices

$$\mathbf{A}_m = \mathbf{A}_{\text{Chain}(d)} + \mathbf{E}_m, \text{ for } m = 3, 4, \dots, 10,$$

where \mathbf{E}_m is the adjacency matrix of the graph $(\bar{V}, \{(1, m)\})$. The graphs we construct are illustrated in Figure 15. Under the null hypothesis, we generate n i.i.d. samples $\mathbf{X}_1, \dots, \mathbf{X}_n \sim N(\mathbf{0}, (\boldsymbol{\Theta}^*(\theta))^{-1})$ where $\boldsymbol{\Theta}^*(\theta) = \mathbf{I}_d + \theta \mathbf{A}_{\text{Chain}(d)}$ and repeat the simulation for $N = 200$ times. Under the alternative, for each repetition, we randomly select $M \sim \text{Unif}(\{3, 4, \dots, 10\})$, set the precision matrix as $\boldsymbol{\Theta}_M^*(\theta) = \mathbf{I}_d + \theta \mathbf{A}_M$ and generate i.i.d. samples $\mathbf{X}_1, \dots, \mathbf{X}_n \sim N(\mathbf{0}, (\boldsymbol{\Theta}_M^*(\theta))^{-1})$. We also repeat this procedure for $N = 200$ times to calculate the type-II error.






Figure 15: $G(\mathbf{A})$ and $G(\mathbf{A}_4)$ for $d = 7$.



(a) $n = 400$



(b) $n = 600$

Figure 16: Type-I + type-II error control for cycles presence test with $\theta \in [0.25, 0.45]$. The grey horizontal line is the significance level $\alpha = 0.05$.

As before the significance level is set to $\alpha = 0.05$, and the tuning parameter as $\lambda = 1.5\sqrt{\log d/n}$. The sample size varies in $n = \{400, 600\}$, the dimension $d = 100, 125$ and 150 and the number of bootstrap samples is $3,000$.

In Figure 16, we plot the risk curves of the test ψ based on Algorithm 4 for $\theta \in [0.25, 0.45]$. With the increase of signal strength or sample size, the risk diminishes to its target value which confirms the results of Section 3.2.3.

Table 2: Sizes for cycle presence testing $\theta \in [0.25, 0.45]$

n	$d \backslash \theta$							
		0.25	0.28	0.32	0.35	0.38	0.42	0.45
400	100	0.010	0.045	0.015	0.045	0.040	0.040	0.055
	125	0.025	0.005	0.045	0.030	0.070	0.070	0.060
	150	0.005	0.005	0.020	0.020	0.015	0.060	0.050
600	100	0.045	0.040	0.060	0.065	0.045	0.060	0.040
	125	0.040	0.040	0.060	0.055	0.065	0.050	0.035
	150	0.030	0.030	0.050	0.015	0.060	0.050	0.040

In Table 2 we report the type-I errors of the tests for each of 6 values of the signal strength θ . The type-I error remains well controlled and is closer to its the target level of 0.05 when the signal strength becomes larger as Corollary 3.4 predicts.

5.3 ADHD-200 Data Analysis

In this section we analyze the difference in connectivity levels of the brain networks of subjects with and without attention deficit hyperactivity disorder (ADHD). We use the ADHD-200 brain-imaging dataset, which is made publicly available by the ADHD-200 consortium (Milham et al., 2012). The dataset consists of 614 subjects with 491 controls and 195 cases.

For each subject a number between 76 and 276 fMRI scans are available. Each of the scans provides a brain imaging vector of 264 voxels, which we use to obtain estimates of the brain network. To simplify the analysis we treat each of the repeated measurements as independent. In addition to the fMRI scans, limited clinical data such as age, gender, handedness etc is also available for each subject. The ages range from 7 to 21 years, with medians 11.75 for the controls and 10.58 for the cases.

Using the ADHD-200 dataset, Cai et al. (2015) explored the connectivity of the brain network, and concluded that the brain networks of ADHD cases are in general less connected compared to the brain networks of healthy subjects. Additionally, Gelfand et al. (2003); Bartzokis et al. (2001) studied how brain connectivity scales with age, and demonstrated that brain connections typically increase from junior to adult age.

Algorithm 8 Multiple Edge Testing

```

Initialize  $\widehat{E}^n \leftarrow \widehat{T}$ ;
repeat
  Reject  $R \leftarrow \{e \in \widehat{E}^n : \sqrt{n}|\widetilde{\Theta}_e| \geq \sqrt{n}\mu + c_{1-\alpha, \widehat{E}^n}\}$ 
  Update  $\widehat{E}^n \leftarrow \widehat{E}^n \setminus R$ 
until  $R = \emptyset$  or  $\widehat{E}^n = \emptyset$ 
return  $\widehat{E}^{\text{nc}} \leftarrow \widehat{T} \setminus \widehat{E}^n$ 

```

To explore potential changes in brain network connectivity in terms of age from junior to senior, we split the dataset into subjects whose age is less than 11 and greater than 11 years, where 11 is chosen since it is close to the median ages of the cases and controls. Moreover, to further study the brain networks at different levels, we define the graph at level $\mu \geq 0$ for the precision matrix Θ^* as

$$G(\Theta^*, \mu) = (\overline{V}, E(\Theta^*, \mu)), \text{ where } E(\Theta^*, \mu) = \{e \mid |\Theta_e^*| \geq \mu\}.$$

When $\mu = 0$, $G(\Theta^*, \mu)$ reduces to the typical graph $G(\Theta^*)$ defined in Section 1.3. We explore the connectivity of $G(\Theta^*, \mu)$ for larger μ 's, because the edges in the brain networks with stronger signal strengths are more informative to reflect the essential topological structure of brains. We therefore consider the hypothesis test

$$\mathbf{H}_0 : G(\Theta^*, \mu) \text{ disconnected vs } \mathbf{H}_1 : G(\Theta^*, \mu) \text{ connected.} \quad (5.3)$$

We can easily modify Algorithm 2 to conduct the test above. In fact, we only need to change the output $\psi_{\alpha, \widehat{T}}^B(\mathcal{D}_2)$ in Algorithm 2 to $\psi_{\alpha, \widehat{T}}^B(\mathcal{D}_2, \mu)$ obtained from Algorithm 8.

Comparing Algorithm 8 with the original Algorithm 1, the only difference is that we change the first step in while loop from $R \leftarrow \{e \in \widehat{E}^n : \sqrt{n}|\widetilde{\Theta}_e| \geq c_{1-\alpha, \widehat{E}^n}\}$ to $R \leftarrow \{e \in \widehat{E}^n : \sqrt{n}|\widetilde{\Theta}_e| \geq \sqrt{n}\mu + c_{1-\alpha, \widehat{E}^n}\}$. Following the same proof as Corollary 3.2, we can show that $\psi_{\alpha, \widehat{T}}^B(\mathcal{D}_2, \mu)$ is a valid test for testing (5.3). A detailed proof can be found in Corollary B.6 in the Appendix. We apply the test $\psi_{\alpha, \widehat{T}}^B(\mathcal{D}_2, \mu)$ to all 4 groups (junior controls, senior controls, junior cases and senior cases) of subjects results for $\mu \in [0.16, 0.3]$. We denote $\widehat{T}(\mu) = \widehat{E}^{\text{nc}}$ where \widehat{E}^{nc} is defined in Algorithm 8. If $\widehat{T}(\mu)$ is connected, since $\widehat{T}(\mu)$ is a tree, we can see that $\psi_{\alpha, \widehat{T}}^B(\mathcal{D}_2, \mu) = 1$ and \mathbf{H}_0 in (5.3) is rejected. For robustness the data is shuffled 100 times before splitting it into two halves and the number of connected components reported is averaged over these 100 data shuffles. Figure 17 demonstrates the number of connected components for $\widehat{T}(\mu)$ when μ varies in the range $[0.16, 0.3]$. We can see that the number of connected components is always larger than 1 and thus we do not reject \mathbf{H}_0 in (5.3) for all $\mu \in [0.16, 0.3]$.

Moreover, we observe in Figure 17 that $\widehat{T}(\mu)$ has more connected components for the controls than the cases. These results confirm the findings of Cai et al. (2015) that the brain

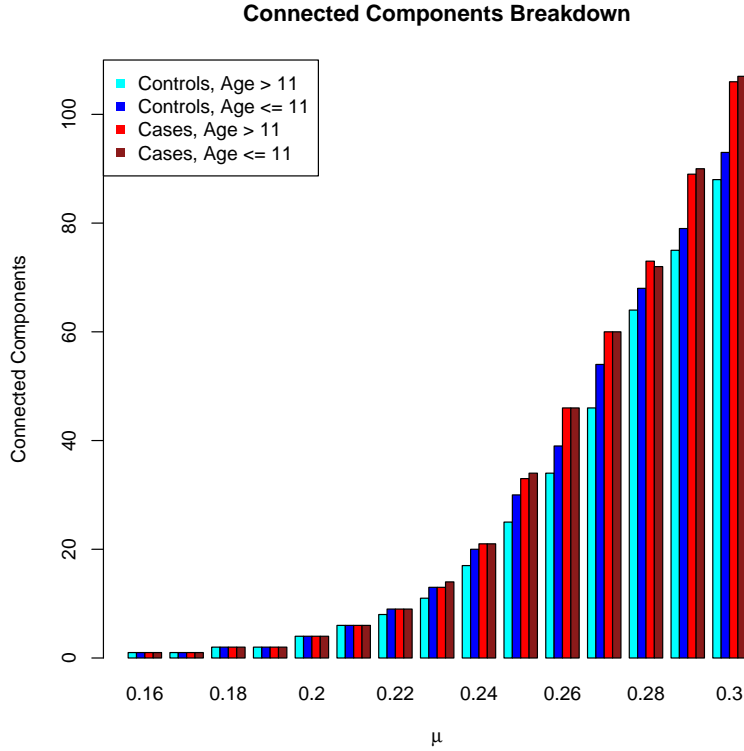


Figure 17: Number of connected components of the graphs when ranging the parameter μ is in the range $[0.16, 0.3]$; We report averaged numbers over 100 data split tests.

network in ADHD patients are less connected compared to the brain network in control subjects at the same age. What is more, Figure 17 also shows that the number of connected components tends to be less in controls whose age ranges within 11 to 21 years as compared to controls who are younger than 11 years of age, which agrees with the results of Gelfand et al. (2003); Bartzokis et al. (2001). Interestingly, there exists less of a difference in the number of components in the brain network of ADHD patients in the two age groups. This observation suggests that subjects with ADHD have a slower rate of forming strong brain connectivity over time. In addition to the the connected components analysis, in Figure 18 we plot the largest connected subgraphs of $\hat{T}(\mu)$ for $\mu = 0.25$ for the controls of age at least 11 years and the top two largest connected subgraphs for the cases. The value $\mu = 0.25$ is chosen as Figure 17 reveals that at $\mu = 0.25$ there exists a clear separation in the number of connected components between cases and controls.

Remarkably, there are two major connected components in the ADHD cases while there is only one major connected component in the controls. In addition, the second major component in the ADHD patients, which is automatically identified by our algorithm, bares close similarity to the salience network (SN) discussed in Cai et al. (2015). Indeed, Cai et al.

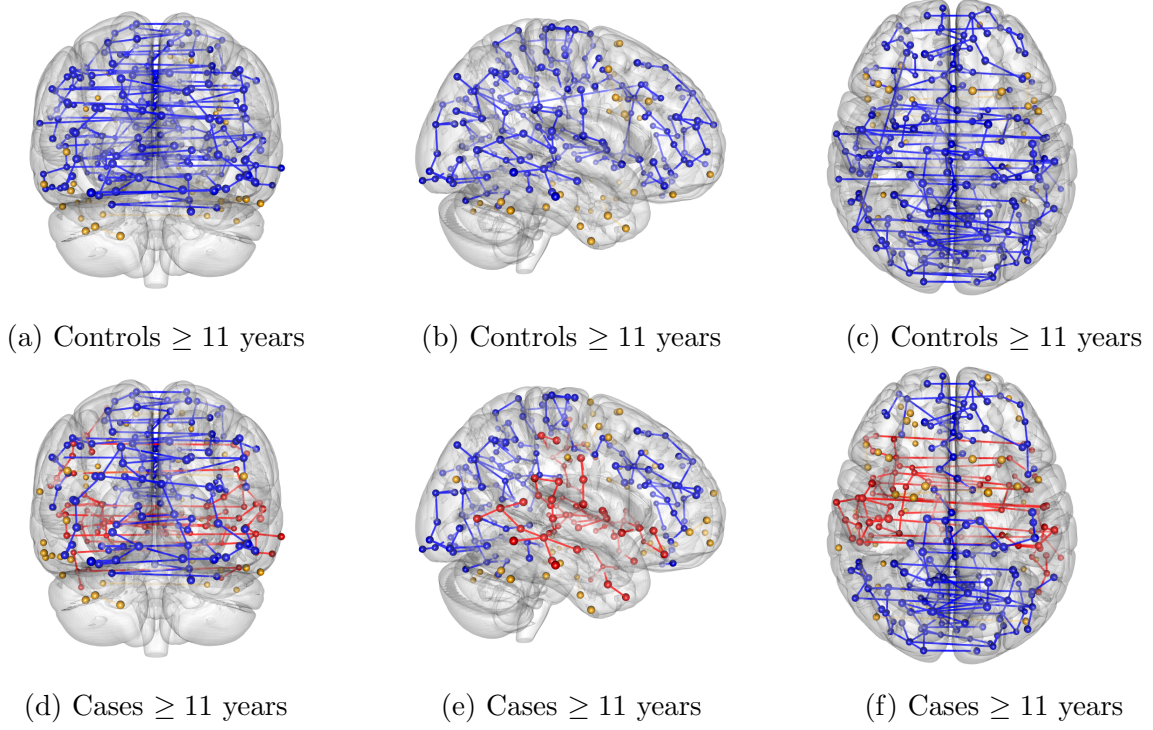


Figure 18: Connected components for age groups ≥ 11 years. Components consisting of at least 6 nodes are displayed in different colors (blue and red), while nodes falling in smaller components are grouped together and plotted in orange.

(2015) studied the connectivity of the SN to two other subgraphs in the brain networks and showed the connectivity between the SN and the other subgraphs is weaker in ADHD cases compared to the connectivity in the controls. This is also shown in Figure 18.

6 Discussion

In the present manuscript we discussed novel strategies for obtaining lower bounds on the signal strength required to perform combinatorial structure tests in graphical models. Whenever single-edge NAS exists, we proved it is necessary for the signal strength to be of order at least $\sqrt{M(\mathcal{C}, d_{G_0}, \log |\mathcal{C}|)/n}$ to ensure successful testing. We further generalized this lower bound strategy to quantify the signal strength needed for testing in cases when multi-edge NAS sets exist, and introduced the concept of buffer entropy. In addition to providing lower bounds, we propose the NAWT, which matches the lower bound limitations on the signal strength and can test combinatorial inference problems under sparsity restrictions. Combinatorial structure tests are potentially applicable in a wide range of disciplines such as epidemiology, computational biology, neurology and economics. We illustrated the last

point via successfully applying a connectivity test to analyze the ADHD-200 dataset. Furthermore we showed that our tests also perform as predicted on synthetic data.

Our lower bound results relate purely combinatorial properties of graphs to signal strength limitations required to perform statistical tests. The bounding techniques are sufficiently general as demonstrated through numerous examples. However, we do recognize that there exists potential room for improvement of our lower bound strategies. One place where modifications leading to sharper results can be introduced is in the construction of precision matrices. Our proofs use matrices based on “one-sided alternatives”, i.e., the signal has the same sign across all non-zero coefficients in the matrix. It is possible that in some cases adding additional uncertainty to the signal through the introduction of a random sign via multiplication by Rademacher random variables, could sharpen the lower bound results. Such an approach was undertaken by [Baraud \(2002\)](#) e.g. Unfortunately, doing so results in a more complicated calculation, and we leave the investigation of such additional randomization devices for future work.

Additionally, the lower bound calculations we performed are tailored for the Gaussian graphical model case. It is of interest to study similar testing limits for other popular models such as the Ising model. Establishing results of this type is not a trivial task given the theory in the present manuscript, and we hope to address these problems in future studies.

A Single-Edge NAS Lower Bound Proofs

Proof of Theorem 2.1. To begin with, we give a high level picture of our proof. The argument consists of three major steps. Our first three steps will show that, there exists a constant R such that if $\theta \leq \frac{1-C^{-1}}{\sqrt{2}(\|\mathbf{A}_0\|_1+2)}$ we have:

$$\gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) \geq 1 - \frac{1}{2} \sqrt{\frac{1}{|\mathcal{C}|^2} \sum_{e, e' \sim \mathcal{C}} \exp \left(n \frac{(R\theta)^{2d_{G_0}(e, e') + 2}}{d_{G_0}(e, e') + 1} \right)} - 1. \quad (\text{A.1})$$

To establish this result, in the first step, we select one precision matrix from the null $\mathcal{S}_0(\theta, s)$ and a set of precision matrices from the alternative $\mathcal{S}_1(\theta, s)$. In the second step, we apply Le Cam's Lemma to the precision matrices constructed above to get a lower bound of $\gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s))$. In the third step, we establish trace perturbation inequalities to further connect the lower bound achieved in the second step to the geometric quantities of the graphs. In the fourth step, we prove the theorem by showing that the right hand side of (A.1) goes to 1 if (2.4) is satisfied.

Step 1 (Matrix Construction).

In this step we construct a class of precision matrices based on the null base graph G_0 and the NAS set \mathcal{C} and verify that these matrices indeed belong to the sets $\mathcal{S}_0(\theta, s)$ and $\mathcal{S}_1(\theta, s)$. We begin with giving the upper bound of matrix norms of adjacent inequalities. Let \mathbf{A}_0 be the adjacency matrix of the graph G_0 . Observe that since \mathbf{A}_0 is symmetric, by Hölder's inequality $\|\mathbf{A}_0\|_2 \leq \sqrt{\|\mathbf{A}_0\|_1 \|\mathbf{A}_0\|_\infty} = \|\mathbf{A}_0\|_1 \leq D$.

Similarly, denote with \mathbf{A}_e the adjacency matrix of the graph $(\bar{V}, \{e\})$ for $e \in \mathcal{C}$. Under our assumptions it follows that $\mathbf{A}_0 + \mathbf{A}_e$ is the adjacency matrix of the graph $G_e = (\bar{V}, E_0 \cup \{e\})$. For brevity, for any two edges $e, e' \in \mathcal{C}$ we define the shorthand notation $\mathbf{A}_{e, e'} := \mathbf{A}_0 + \mathbf{A}_e + \mathbf{A}_{e'}$. Take $\Theta_0 = \mathbf{I} + \theta \mathbf{A}_0$, $\Theta_e = \mathbf{I} + \theta(\mathbf{A}_0 + \mathbf{A}_e)$, $\Theta_{e, e'} = \mathbf{I} + \theta \mathbf{A}_{e, e'}$, for $e, e' \in \mathcal{C}$, $\theta > 0$. By the triangle inequality for any $e, e' \in \mathcal{C}$ we have:

$$\begin{aligned} \max(\|\mathbf{A}_0\|_2, \|\mathbf{A}_0 + \mathbf{A}_e\|_2, \|\mathbf{A}_{e, e'}\|_2) &\leq \|\mathbf{A}_0\|_2 + 2, \\ \max(\|\mathbf{A}_0\|_1, \|\mathbf{A}_0 + \mathbf{A}_e\|_1, \|\mathbf{A}_{e, e'}\|_1) &\leq \|\mathbf{A}_0\|_1 + 2. \end{aligned}$$

Next we make sure that the matrices Θ_0 and Θ_e fall into the set $\mathcal{M}(s)$ and in addition the matrix $\Theta_{e, e'} > 0$. For the upper bounds, it suffices to choose η satisfying:

$$\max(\|\Theta_0\|_2, \|\Theta_e\|_2) \leq 1 + (\|\mathbf{A}_0\|_2 + 2)\theta \leq C, \quad \max(\|\Theta_0\|_1, \|\Theta_e\|_1) \leq 1 + (\|\mathbf{A}_0\|_1 + 2)\theta \leq L.$$

Recall that $\|\mathbf{A}_0\|_2 \leq \|\mathbf{A}_0\|_1$, and $C \leq L$ hence both inequalities are implied if $1 + (\|\mathbf{A}_0\|_1 + 2)\theta \leq C$. Furthermore, by Weyl's inequality:

$$\lambda_d(\Theta_0), \lambda_d(\Theta_e), \lambda_d(\Theta_{e, e'}) \geq 1 - \theta(\|\mathbf{A}_0\|_2 + 2) \geq 1 - \theta(\|\mathbf{A}_0\|_1 + 2), \quad (\text{A.2})$$

where λ_d denotes the smallest eigenvalue of the corresponding matrix. We want to ensure that the last term is $\geq C^{-1}$. Since by assumption $\theta < \frac{1-C^{-1}}{\sqrt{2(D+2)}} = \frac{1-C^{-1}}{\sqrt{2(\|\mathbf{A}_0\|_1+2)}}$ the above inequalities are satisfied (note that trivially $1 - C^{-1} \leq C - 1$ since $C \geq 1$). We also obtain as a by-product that $\Theta_{e,e'} \geq 0$.

Step 2 (*Risk Lower Bound*).

In this step we obtain a lower bound on the minimax risk driven by Le Cam's Lemma (Le Cam, 1973). Using a determinant identity we control the chi-divergence by the traces of adjacency matrices's powers. By assumption $G_0 \in \mathcal{G}_0, G_e \in \mathcal{G}_1$ for all $e \in \mathcal{C}$ and hence the induced graphs $G(\Theta_0) \in \mathcal{G}_0$ and $G(\Theta_e) \in \mathcal{G}_1$ for all $e \in \mathcal{C}$. Put the uniform prior on \mathcal{C} and consider the models generated by $N(0, (\Theta_e)^{-1})$ where $e \in \mathcal{C}$. Define:

$$\bar{\mathbb{P}} = \frac{1}{|\mathcal{C}|} \sum_{e \in \mathcal{C}} \mathbb{P}_{\Theta_e},$$

where \mathbb{P}_{Θ_e} we define the probability measure when the data is i.i.d. $\mathbf{X}_i \sim N(0, (\Theta_e)^{-1})$, and let \mathbb{P}_{Θ_0} be the probability measure when the data is i.i.d. $\mathbf{X}_i \sim N(0, (\Theta_0)^{-1})$. Next, by Neyman-Pearson's lemma we have:

$$\gamma(\mathcal{S}_0, \mathcal{S}_1) \geq \inf_{\psi} \left[\mathbb{P}_{\Theta_0}(\psi = 1) + \bar{\mathbb{P}}(\psi = 0) \right] = 1 - \text{TV}(\bar{\mathbb{P}}, \mathbb{P}_{\Theta_0}), \quad (\text{A.3})$$

where for two probability measures $P, Q \ll \lambda$ on a measurable space (Ω, \mathcal{A}) , TV stands for *total variation distance*, and is defined as

$$\text{TV}(P, Q) = \sup_{A \in \mathcal{A}} |P(A) - Q(A)| = \frac{1}{2} \int \left| \frac{dP}{d\lambda}(\omega) - \frac{dQ}{d\lambda}(\omega) \right| d\lambda(\omega).$$

By Cauchy-Schwartz one has:

$$1 - \text{TV}(\bar{\mathbb{P}}, \mathbb{P}_{\Theta_0}) \geq 1 - \frac{1}{2} \sqrt{D_{\chi^2}(\bar{\mathbb{P}}, \mathbb{P}_{\Theta_0})}, \quad (\text{A.4})$$

where $D_{\chi^2}(P, Q)$ is the chi-square divergence between the measures P, Q and is defined as:

$$D_{\chi^2}(P, Q) = \int \left(\frac{dP}{dQ}(\omega) - 1 \right)^2 dQ(\omega) = \int \left(\frac{dP}{dQ}(\omega) \right)^2 dQ(\omega) - 1,$$

assuming that $P \ll Q$. Observe that $D_{\chi^2}(\bar{\mathbb{P}}, \mathbb{P}_{\Theta_0})$ can be equivalently expressed as:

$$D_{\chi^2}(\bar{\mathbb{P}}, \mathbb{P}_{\Theta_0}) = \mathbb{E}_{\Theta_0} L_{\Theta_0}^2 - 1, \quad (\text{A.5})$$

where $L_{\Theta_0} = \frac{1}{|\mathcal{C}|} \sum_{e \in \mathcal{C}} \frac{d\mathbb{P}_{\Theta_e}}{d\mathbb{P}_{\Theta_0}}$ is the integrated likelihood ratio, and \mathbb{E}_{Θ_0} denotes the expectation under $\mathbf{X}_i \sim N(0, (\Theta_0)^{-1})$. Hence by (A.3) and (A.4), it suffices to obtain upper bounds on the integrated likelihood ratio in order to lower bound the minimax risk (1.3). Writing

out the likelihood ratio comparing the normal distribution with precision matrix Θ_0 to the uniform mixture of normal distribution with precision matrix Θ_e for $e \in \mathcal{C}$ we get:

$$L_{\Theta_0} = \frac{1}{|\mathcal{C}|} \sum_{e \in \mathcal{C}} \left(\frac{\det(\Theta_e)}{\det(\Theta_0)} \right)^{n/2} \prod_{i=1}^n \exp(-\mathbf{X}_i^T \theta \mathbf{A}_e \mathbf{X}_i / 2).$$

To calculate the chi-square distance in (A.5), next we square this expression and take its expectation under \mathbb{P}_{Θ_0} to obtain:

$$\begin{aligned} \mathbb{E}_{\Theta_0} L_{\Theta_0}^2 &= \frac{1}{|\mathcal{C}|^2} \sum_{e, e' \in \mathcal{C}} \left(\frac{\det(\Theta_e)}{\det(\Theta_0)} \right)^{n/2} \left(\frac{\det(\Theta_{e'})}{\det(\Theta_0)} \right)^{n/2} \mathbb{E}_{\Theta_0} \left[\prod_{i=1}^n \exp(-\mathbf{X}_i^T \theta (\mathbf{A}_e + \mathbf{A}_{e'}) \mathbf{X}_i / 2) \right] \\ &= \frac{1}{|\mathcal{C}|^2} \sum_{e, e' \in \mathcal{C}} \left(\frac{\det(\Theta_e)}{\det(\Theta_0)} \right)^{n/2} \left(\frac{\det(\Theta_{e'})}{\det(\Theta_{e, e'})} \right)^{n/2}. \end{aligned} \quad (\text{A.6})$$

Next, we will expand the determinants above. Recall that we have ensured that $1 - \theta(\|\mathbf{A}_0\|_2 + 2) > 0$ (see A.2). This implies

$$\theta \max(\|\mathbf{A}_0\|_2, \|\mathbf{A}_0 + \mathbf{A}_e\|_2, \|\mathbf{A}_0 + \mathbf{A}_{e'}\|_2, \|\mathbf{A}_{e, e'}\|_2) \leq 1.$$

For what follows for a symmetric matrix $\mathbf{A}_{d \times d}$ we denote its ordered eigenvalues with $\lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \dots \geq \lambda_d(\mathbf{A})$. Let $\mathbf{A} \in \mathbb{R}^{d \times d}$ be a symmetric matrix such that $\|\mathbf{A}\|_2 \leq 1$. Then we have:

$$\begin{aligned} \log \det(\mathbf{I} + \mathbf{A}) &= \sum_{j=1}^d \log \lambda_j(\mathbf{I} + \mathbf{A}) = \sum_{j=1}^d \log(1 + \lambda_j(\mathbf{A})) = \sum_{j=1}^d \sum_{k=1}^{\infty} (-1)^{k+1} \frac{\lambda_j^k(\mathbf{A})}{k} \\ &= \sum_{k=1}^{\infty} (-1)^{k+1} \frac{\text{Tr}(\mathbf{A}^k)}{k}. \end{aligned}$$

Using the form $\det(\mathbf{I} + \mathbf{A}) = \exp(\log \det(\mathbf{I} + \mathbf{A}))$ and plugging the above equation into (A.6), we conclude that:

$$\mathbb{E}_{\Theta_0} L_{\Theta_0}^2 = \frac{1}{|\mathcal{C}|^2} \sum_{e, e' \in \mathcal{C}} \exp \left(\frac{n}{2} \sum_{k=1}^{\infty} \frac{(-\theta)^k}{k} (T_1^k + T_2^k) \right), \text{ where}$$

$$T_1^k := \text{Tr}[\mathbf{A}_0^k - (\mathbf{A}_0 + \mathbf{A}_e)^k] \quad T_2^k := \text{Tr}[(\mathbf{A}_{e, e'})^k - (\mathbf{A}_0 + \mathbf{A}_{e'})^k].$$

Step 3 (*Trace Perturbation Inequalities*).

In this step, we control $T_1^k + T_2^k$ in terms of k and link it with the geometric quantities of the graph. We view T_1^k as the perturbation difference between $\text{Tr}[\mathbf{A}_0^k]$ and $\text{Tr}[(\mathbf{A}_0 + \mathbf{A}_e)^k]$ and

we treat T_2^k similarly. In the following step, we aim to develop the perturbation inequalities for the trace of matrix powers.

First we will argue that $T_1^k + T_2^k \geq 0$ for all even $k \in \mathbb{N}$. To see this recall that the trace operator of an adjacency matrix \mathbf{M} satisfies

$$\text{Tr}(\mathbf{M}^k) = \text{number of all closed walks of length } k.$$

First we consider case $e \neq e'$. Notice that all closed walks in $G(\mathbf{A}_0 + \mathbf{A}_e)$ that do not belong to $G(\mathbf{A}_0)$ have to pass through the edge e at least once. Similarly all closed walks in $G(\mathbf{A}_0 + \mathbf{A}_{e'})$ that do not belong to $G(\mathbf{A}_0)$ have to pass through the edge e' at least once. Furthermore, all closed walks of length k passing through either e or e' belong to $G(\mathbf{A}_{e,e'})$. In addition $G(\mathbf{A}_{e,e'})$ might contain extra closed walks passing through both e and e' . This shows:

$$T_1^k + T_2^k \geq 0,$$

for all k . Since when k is odd we have $(-\theta)^k(T_1^k + T_2^k) \leq 0$, we only need to focus on k even.

Next we prove that for $k < 2d_{G_0}(e, e') + 2$, we have $T_1^k + T_2^k \equiv 0$. To see this, first consider the case $e \neq e'$. Notice that the graph $G(\mathbf{A}_{e,e'})$ cannot contain paths passing through both e and e' unless $k \geq 2d_{G_0}(e, e') + 2$. To see this, notice that no even length closed walk between e and e' can exist if the length of this walk is smaller than $2d_{G_0}(e, e')$ plus the two edges e and e' . This proves our claim in the case $e \neq e'$. In the special case $e = e'$, the length of the path trivially needs to be at least of length 2 to pass through both e and e' .

We will now argue that for even $k \in \mathbb{N}$ we have $T_1^k + T_2^k \leq 2(\|\mathbf{A}_0\|_2 + 2)^k$. In fact we will prove that $T_1^k \leq 0$ for all k and $T_2^k \leq 2(\|\mathbf{A}_0\|_2 + 2)^k$ for all even k . To see that $T_1^k \leq 0$, note that $G(\mathbf{A}_0)$ contains less closed walks than $G(\mathbf{A}_0 + \mathbf{A}_e)$.

Recall that for a symmetric matrix $\mathbf{A}_{d \times d}$ we denote its ordered eigenvalues with $\lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \dots \geq \lambda_d(\mathbf{A})$. Using Corollary D.2 for the matrices $\mathbf{A}_{e,e'}$, $\mathbf{A}_0 + \mathbf{A}_{e'}$, \mathbf{A}_e with constants

$$c_{\sigma(j)} = \text{sign}(\lambda_j(\mathbf{A}_{e,e'}) - \lambda_j(\mathbf{A}_e)) |\lambda_j(\mathbf{A}_{e,e'}) - \lambda_j(\mathbf{A}_e)|^{k-1},$$

we obtain:

$$\sum_{j=1}^d |\lambda_j(\mathbf{A}_{e,e'}) - \lambda_j(\mathbf{A}_e)|^k \leq \sum_{j=1}^d c_j \lambda_j(\mathbf{A}_0 + \mathbf{A}_{e'}) \leq \left[\sum_{j=1}^d |c_j|^{\frac{k}{k-1}} \right]^{\frac{k-1}{k}} \left[\sum_{j=1}^d |\lambda_j(\mathbf{A}_0 + \mathbf{A}_{e'})|^k \right]^{\frac{1}{k}},$$

where the last inequality follows by Hölder's inequality. We conclude that:

$$\sum_{j=1}^d |\lambda_j(\mathbf{A}_{e,e'}) - \lambda_j(\mathbf{A}_e)|^k \leq \sum_{j=1}^d |\lambda_j(\mathbf{A}_0 + \mathbf{A}_{e'})|^k. \quad (\text{A.7})$$

Next, observe that for a single edge matrices the matrix $-\mathbf{A}_e$ has very simple eigenvalue structure: 1, -1 and $d-2$ zeros. Hence we conclude that for even k :

$$\begin{aligned} T_2^k &= \text{Tr}((\mathbf{A}_{e,e'})^k) - \text{Tr}((\mathbf{A}_0 + \mathbf{A}_{e'})^k) = \sum_{j=1}^d |\lambda_j(\mathbf{A}_{e,e'})|^k - \sum_{j=1}^d |\lambda_j(\mathbf{A}_0 + \mathbf{A}_{e'})|^k \\ &\leq |\lambda_1(\mathbf{A}_{e,e'})|^k - |(\lambda_1(\mathbf{A}_{e,e'}) - 1)|^k + |\lambda_d(\mathbf{A}_{e,e'})|^k - |\lambda_d(\mathbf{A}_{e,e'}) + 1|^k \end{aligned} \quad (\text{A.8})$$

$$\begin{aligned} &\leq |\lambda_1(\mathbf{A}_{e,e'})|^k + |\lambda_d(\mathbf{A}_{e,e'})|^k \\ &\leq 2\|\mathbf{A}_{e,e'}\|_2^k \leq 2(\|\mathbf{A}_0\|_2 + 2)^k. \end{aligned} \quad (\text{A.9})$$

The last shows that indeed $T_1^k + T_2^k \leq 2(\|\mathbf{A}_0\|_2 + 2)^2$ as claimed. Putting everything together we obtain

$$\begin{aligned} \sum_{k=1}^{\infty} \frac{(-\theta)^k}{k} [T_1^k + T_2^k] &\leq \sum_{2|k, k \geq 2d_{G_0}(e,e')+2}^{\infty} \frac{\theta^k}{k} [T_1^k + T_2^k] \\ &\leq \sum_{2|k, k \geq 2d_{G_0}(e,e')+2}^{\infty} \frac{2((\|\mathbf{A}_0\|_2 + 2)\theta)^k}{k} \leq \frac{2((\|\mathbf{A}_0\|_2 + 2)\theta)^{2d_{G_0}(e,e')+2}}{(2d_{G_0}(e,e') + 2)(1 - (\theta(\|\mathbf{A}_0\|_2 + 2))^2)} \\ &\leq \frac{2((\|\mathbf{A}_0\|_2 + 2)\theta)^{2d_{G_0}(e,e')+2}}{d_{G_0}(e,e') + 1}, \end{aligned}$$

where in the last inequality we used the fact that $\theta \leq \frac{1}{\sqrt{2}(\|\mathbf{A}_0\|_2 + 2)}$ which follows by the requirements on θ . This completes the proof of (A.1) where $R = \sqrt{2}(\|\mathbf{A}_0\|_2 + 2)$.

Step 4 (Rate Control).

The goal in this final step is to show that if (2.4) holds the risk $\liminf \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1$. The proof is technical, but the high-level idea is to clip the first $\log |\mathcal{C}|$ degrees in (A.1) and deal with two separate summations. It turns out that the scaling assumed on θ in (2.4) is precisely enough to control both the summation of all degrees below $\log |\mathcal{C}|$ and the summation of all degrees above $\log |\mathcal{C}|$. Define the following quantities:

$$K_r := |\{(e, e') : e, e' \in \mathcal{C}, d_{G_0}(e, e') = r\}|,$$

where (e, e') are unordered edge pairs, and observe that $\sum_{r \geq 0} K_r = \binom{|\mathcal{C}|}{2} + |\mathcal{C}|$ by definition. We will in fact, first show that if $\theta \leq \kappa \sqrt{\frac{\log |\mathcal{C}|}{n}}$ for some small κ , and

$$\sum_{r=0}^{\lfloor \log |\mathcal{C}| \rfloor} K_r = O(|\mathcal{C}|^{2-\gamma}), \quad (\text{A.10})$$

for some $0 < \gamma \leq 1$, then $\limsup \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1$ provided that $|\mathcal{C}| \rightarrow \infty$. We will then derive the Theorem as a Corollary to this observation. On an important note, in Remark

A.1 we will argue that if condition (A.10), is satisfied then $M(\mathcal{C}, d_{G_0}, \frac{1}{2} \log |\mathcal{C}|) \asymp \log |\mathcal{C}|$, and hence we do not lose any generality in drawing the conclusion of the Theorem as a Corollary to this setup.

It suffices to control:

$$\underbrace{\frac{2}{|\mathcal{C}|^2} \sum_{(e,e'): d_{G_0}(e,e') \geq 1} \exp\left(n \frac{\bar{\theta}^{2d_{G_0}(e,e')+2}}{d_{G_0}(e,e') + 1}\right)}_{I_1} + \underbrace{\frac{2K_0 - |\mathcal{C}|}{|\mathcal{C}|^2} \exp(n\bar{\theta}^2/2)}_{I_2}, \quad (\text{A.11})$$

where we will write $\bar{\theta}$ for $R\theta$ for brevity. We will show that when $\bar{\theta}$ is small, (A.11) is bounded by 1 asymptotically, which in turn suffices to show that $\limsup \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1$. Notice that $\bar{\theta}$ and θ are the same quantity up to the constant R and hence $\theta < \kappa \sqrt{\frac{\log |\mathcal{C}|}{n}}$ is equivalent to $\bar{\theta} < \bar{\kappa} \sqrt{\frac{\log |\mathcal{C}|}{n}}$ for some sufficiently small $\bar{\kappa}$. We will require:

$$\bar{\kappa} < \sqrt{2\gamma} \wedge (ec_0)^{-1/2}.$$

Observe that since $K_0 = O(|\mathcal{C}|^{2-\gamma})$, $\bar{\kappa}^2/2 < \gamma$, and $\bar{\theta} < \bar{\kappa} \sqrt{\frac{\log |\mathcal{C}|}{n}}$ we have:

$$I_2 \leq \frac{(2K_0 - |\mathcal{C}|)|\mathcal{C}|^{\bar{\kappa}^2/2}}{|\mathcal{C}|^2} \rightarrow 0.$$

Next we tackle the term I_1 in (A.11). We will show that since $\bar{\theta} = \bar{\kappa} \sqrt{\frac{\log |\mathcal{C}|}{n}} = o(1)$ by assumption, this term goes to 1.

$$\begin{aligned} I_1 &= \frac{2}{|\mathcal{C}|^2} \sum_{r \geq 1} K_r \exp(n\bar{\theta}^{2r+2}/(r+1)) = \frac{2}{|\mathcal{C}|^2} \sum_{r=1}^{d-1} K_r |\mathcal{C}|^{\bar{\kappa}^2 \bar{\theta}^{2r}/(r+1)} + \frac{2K_\infty}{|\mathcal{C}|^2} \\ &< \frac{2}{|\mathcal{C}|^2} \sum_{r=1}^{d-1} K_r |\mathcal{C}|^{2\bar{\theta}^{2r}/(r+1)} + \frac{2K_\infty}{|\mathcal{C}|^2}, \end{aligned}$$

where the last inequality follows by the fact that $\bar{\kappa}^2/2 < \gamma < 1$. Splitting out the first $\lfloor \log |\mathcal{C}| \rfloor$ terms out of this summation yields:

$$I_1 < \underbrace{\frac{2}{|\mathcal{C}|^2} \sum_{r=1}^{\lfloor \log |\mathcal{C}| \rfloor} K_r |\mathcal{C}|^{2\bar{\theta}^{2r}/(r+1)}}_{I_{11}} + \underbrace{\frac{2}{|\mathcal{C}|^2} \sum_{r=\lfloor \log |\mathcal{C}| \rfloor + 1}^{d-1} K_r |\mathcal{C}|^{2\bar{\theta}^{2r}/(r+1)} + \frac{2K_\infty}{|\mathcal{C}|^2}}_{I_{12}}.$$

The first term is bounded by:

$$I_{11} \leq 2 \left(\sum_{r=1}^{\lfloor \log |\mathcal{C}| \rfloor} K_r \right) \frac{|\mathcal{C}|^{\bar{\theta}^2}}{|\mathcal{C}|^2} = o(1),$$

where we used $\left(\sum_{r=2}^{\lfloor \log |\mathcal{C}| \rfloor} K_r\right) = O(|\mathcal{C}|^{2-\gamma})$ and the fact that $\bar{\theta}^2 < \gamma$. Next, notice that

$$|\mathcal{C}|^{2\bar{\theta}^{2r}/(r+1)} \leq 1 + 6\bar{\theta}^{2r}(|\mathcal{C}| - 1)/(r + 1).$$

This follows by:

$$\exp(2\log(|\mathcal{C}|)\bar{\theta}^{2r}/(r + 1)) \leq 1 + 6\log(|\mathcal{C}|)\bar{\theta}^{2r}/(r + 1) \leq 1 + 6(|\mathcal{C}| - 1)\bar{\theta}^{2r}/(r + 1),$$

with the first inequality holding when $2\log(|\mathcal{C}|)\bar{\theta}^{2r}/(r + 1) < 1$, which is true since $\bar{\theta} < \bar{\kappa}\sqrt{c_0} < \exp(-1/2) < \frac{1}{\sqrt{2}}$, and $r \geq \lfloor \log |\mathcal{C}| \rfloor + 1$. Hence we have:

$$\begin{aligned} I_{12} &\leq \frac{2}{|\mathcal{C}|^2} \sum_{r=\lfloor \log |\mathcal{C}| \rfloor + 1}^{d-1} K_r (1 + 6(|\mathcal{C}| - 1)\bar{\theta}^{2r}/(r + 1)) + \frac{2K_\infty}{|\mathcal{C}|^2} \\ &\leq \left(1 - \frac{O(|\mathcal{C}|^{2-\gamma})}{|\mathcal{C}|^2}\right) + \frac{12(|\mathcal{C}| - 1)}{|\mathcal{C}|^2} \sum_{r=\lfloor \log |\mathcal{C}| \rfloor + 1}^{d-1} \frac{K_r \bar{\theta}^{2r}}{r} \\ &\leq \left(1 - \frac{O(|\mathcal{C}|^{2-\gamma})}{|\mathcal{C}|^2}\right) + \frac{12\bar{\theta}^{2\lfloor \log |\mathcal{C}| \rfloor + 2}}{1 - \bar{\theta}^2} |\mathcal{C}|^{-1} \max_{\lfloor \log |\mathcal{C}| \rfloor + 1 \leq r} \frac{K_r}{r}. \end{aligned}$$

Paying closer attention to the second term we have:

$$\begin{aligned} \frac{12\bar{\theta}^{2\lfloor \log |\mathcal{C}| \rfloor + 2}}{1 - \bar{\theta}^2} |\mathcal{C}|^{-1} \max_{\lfloor \log |\mathcal{C}| \rfloor + 1 \leq r} \frac{K_r}{r} &\leq 24\bar{\theta}^{2\lfloor \log |\mathcal{C}| \rfloor + 2} \frac{\binom{|\mathcal{C}|}{2} + |\mathcal{C}|}{|\mathcal{C}|} \\ &\leq 24\bar{\theta}^{2\lfloor \log |\mathcal{C}| \rfloor + 2} |\mathcal{C}| = o(1), \end{aligned}$$

with the last equality holds since:

$$\bar{\theta} \leq \bar{\kappa}\sqrt{c_0} < \exp(-1/2),$$

as we required. This combined with (A.1) concludes the proof of $\limsup \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) = 1$ when $\theta \leq \kappa\sqrt{\frac{\log |\mathcal{C}|}{n}}$.

Finally, notice that any subset of a NAS \mathcal{C} is trivially a NAS. Hence we can apply what we just showed to the set $N_{\log |\mathcal{C}|} \subset \mathcal{C}$ — the maximal $\log |\mathcal{C}|$ -packing of \mathcal{C} . Evaluating the constants K_r on $N_{\log |\mathcal{C}|}$ gives:

$$K_0 = |N_{\log |\mathcal{C}|}|, K_r = 0 \text{ for all } r \leq \lfloor \log |\mathcal{C}| \rfloor,$$

and since $\lfloor N_{\log |\mathcal{C}|} \rfloor \leq \lfloor \log |\mathcal{C}| \rfloor$ we conclude that $\sum_{r=0}^{\lfloor N_{\log |\mathcal{C}|} \rfloor} K_r = |N_{\log |\mathcal{C}|}| = O(|N_{\log |\mathcal{C}|}|^{2-\gamma})$ for any $0 < \gamma \leq 1$. \square

Remark A.1. In this remark we argue that if one can find a NAS \mathcal{C} with null base G_0 satisfying $|\mathcal{C}| \rightarrow \infty$ and (A.10), then $M(\mathcal{C}, d_{G_0}, \frac{1}{2} \log |\mathcal{C}|)$ satisfies $M(\mathcal{C}, d_{G_0}, \frac{1}{2} \log |\mathcal{C}|) \asymp \log |\mathcal{C}|$, and hence there are no benefits on imposing conditions on K_r as in (A.10) over conditions on the entropy except in terms of constants. We build a set $N_{\log |\mathcal{C}|/2}$ in a sequential manner, by transferring an edge e from \mathcal{C} and removing all edges $e' \in \mathcal{C}$ such that $d_{G_0}(e, e') \leq \frac{1}{2} \log |\mathcal{C}|$. Suppose in total there were k such additions so that $\log k \leq M(\mathcal{C}, d_{G_0}, \frac{1}{2} \log |\mathcal{C}|)$. If on the i -th step there were m_i edges removed then by the triangle inequality and (A.10) we have:

$$\sum_{i \in [k]} \binom{m_i}{2} + m_i = O(|\mathcal{C}|^{2-\gamma}),$$

and thus $\sum_{i \in [k]} m_i^2 = O(|\mathcal{C}|^{2-\gamma})$. Hence since we have removed all edges from $|\mathcal{C}|$ we conclude:

$$\frac{|\mathcal{C}|}{k} \leq \frac{\sum_{i \in [k]} m_i}{k} \leq \sqrt{\frac{\sum_{i \in [k]} m_i^2}{k}} = O\left(\frac{|\mathcal{C}|^{1-\gamma/2}}{\sqrt{k}}\right),$$

which shows that $k \gtrsim |\mathcal{C}|^\gamma$, and therefore indeed $M(\mathcal{C}, d_{G_0}, \frac{1}{2} \log |\mathcal{C}|) \asymp \log |\mathcal{C}|$. Finally we note that a constant in front of the the packing radius $\log |\mathcal{C}|$ will not affect the rates. This completes the proof.

Proof of Theorem 2.2. The proof of this result follows the same line of argument as the proof of Theorem 2.1. Here we just sketch the differences. The first important observation is that the risk $\gamma(\mathcal{S}_0, \mathcal{S}_1)$ is symmetric in the sense that $\gamma(\mathcal{S}_0, \mathcal{S}_1) = \gamma(\mathcal{S}_1, \mathcal{S}_0)$. This implies that if controlling the eigenvalues of precision matrices made in the proof of Theorem 2.1 hold for edge deletion NAS, the proof will continue to hold. If we denote by \mathbf{A}_e the adjacency matrix of $(\bar{V}, \{e\})$ for any $e \in \mathcal{C}$, then the adjacency matrices of the graphs under the null $G_{\setminus e}$ are given by $\mathbf{A}_1 - \mathbf{A}_e$. It is now a simple exercise to check that:

$$\text{Tr}(\mathbf{A}_1^k) - \text{Tr}((\mathbf{A}_1 - \mathbf{A}_e)^k) \leq 2\|\mathbf{A}_1\|_2^k, \quad \text{Tr}((\mathbf{A}_1 - \mathbf{A}_{e'})^k) - \text{Tr}((\mathbf{A}_1 - \mathbf{A}_e - \mathbf{A}_{e'})^k) \leq 0,$$

and the proof is completed as in Theorem 2.1. □

B Structure Testing Algorithms Proofs

B.1 Connectivity Testing Proofs

Lemma B.1. Assume the estimate $\hat{\Theta}^{(1)}$ satisfies (3.3) on the parameter space $\mathcal{M}(s)$ and let $\theta = rK\sqrt{\log d/\lfloor n/2 \rfloor}$, where $r > 2$. Then for all $e \in E(\hat{T})$, we have $|\Theta_e^*| \geq \frac{r-2}{r}\theta$ when $\Theta^* \in \mathcal{S}_1(\theta, s)$ (see definition in (3.12)) with probability at least $1 - 1/d$. Moreover, when $\Theta^* \in \mathcal{S}_0(s)$ (see definition in (3.11)) for $\hat{E} = \text{argmin}_{e \in E(\hat{T})} |\Theta_e^{(1)}|$ we have $\Theta_{\hat{E}}^* = 0$ with probability at least $1 - 1/d$.

Proof of Lemma B.1. By (3.3) we obtain the following:

$$|\widehat{\Theta}_e^{(1)}| \geq \frac{r-1}{r}\theta, \text{ when } |\Theta_e^*| \geq \theta, \quad (\text{B.1})$$

$$|\widehat{\Theta}_e^{(1)}| \leq \frac{\theta}{r}, \text{ when } |\Theta_e^*| = 0, \quad (\text{B.2})$$

with probability at least $1 - 1/d$ uniformly for all e . First consider the case when $\Theta^* \in \mathcal{S}_0(s)$. Since the graph is disconnected, we have that there exists at least one edge in $e \in \widehat{T}$, such that $\Theta_e^* = 0$ and hence by (B.2) we have $|\widehat{\Theta}_e^{(1)}| \leq \theta/r$. This implies that: $|\widehat{\Theta}_{\widehat{E}}^{(1)}| \leq \theta/r$. Now suppose that $\Theta_{\widehat{E}}^* \neq 0$. Denote the connected component of $G(\Theta^*)$ containing \widehat{E} with \mathcal{C}^* . Since $\Theta^* \in \mathcal{S}_0(s)$ there exists a spanning tree $T_{\mathcal{C}^*}$ of \mathcal{C}^* such that for all $e \in E(T_{\mathcal{C}^*})$ we have $|\Theta_e^*| \geq \theta$. By (B.1) the latter implies that $|\widehat{\Theta}_e^{(1)}| > |\widehat{\Theta}_{\widehat{E}}^{(1)}|$ for all $e \in E(T_{\mathcal{C}^*})$. On the other hand, note that the value $|\widehat{\Theta}_{\widehat{E}}^{(1)}|$ is the largest among the values of all edges connecting the two connected components of $G(\widehat{\Theta}^{(1)})$ which \widehat{E} connects. However, as we argued there exist an edge $\tilde{e} \in T_{\mathcal{C}^*}$ connecting these two components satisfying $|\widehat{\Theta}_{\tilde{e}}^{(1)}| > |\widehat{\Theta}_{\widehat{E}}^{(1)}|$. This is a contradiction with the fact that \widehat{T} is a MST. Hence $\Theta_{\widehat{E}}^* = 0$.

Next let $\Theta^* \in \mathcal{S}_1(\theta, s)$. In this case by (B.1) the tree will cover only edges with $|\widehat{\Theta}_e| \geq (r-1)\theta/r$ and hence by assumption (3.3) and the fact that $\Theta^* \in \mathcal{S}_1(\theta, s)$ we have $|\Theta_{\widehat{E}}^*| \geq (r-2)\theta/r$. This completes the proof. \square

Proof of Corollary 3.2. We first verify the size of the test. By Lemma B.1 we conclude that with asymptotic probability 1 we have $\min_{e \in E(\widehat{T})} |\Theta_e^*| = 0$ when $\Theta^* \in \mathcal{S}_0(s)$. Conditioning on this event by Proposition 3.1 we conclude that:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(\theta)} \mathbb{P}(\psi = 1) \leq \alpha.$$

For the second part, with asymptotic probability 1 we have that $|\Theta_{\widehat{E}}^*| \geq \theta \frac{\epsilon}{2+\epsilon} \geq \frac{\sqrt{2}\epsilon}{2+\epsilon} K \sqrt{\log d/n}$. Using the second part of Proposition 3.1 we can guarantee:

$$\liminf_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\psi = 0) = 0.$$

This completes the proof. \square

Lemma B.2. Assume the estimate $\widehat{\Theta}^{(1)}$ satisfies assumption (3.3) on the parameter space $\mathcal{M}(s)$ and $\theta \geq rK\sqrt{\log d/\lfloor n/2 \rfloor}$ for some $r > 2$. If $\Theta^* \in \mathcal{S}_0(s)$ then $|\Theta_{\widehat{E}}^*| = 0$ where $\widehat{E} = \operatorname{argmin}_{e \in E(\widehat{F})} |\widehat{\Theta}_e^{(1)}|$ with probability at least $1 - d^{-1}$. Furthermore, when $\Theta^* \in \mathcal{S}_1(\theta, s)$, for all $e \in E(\widehat{F})$ we have $|\Theta_e^*| \geq \frac{r-2}{r}\theta$ with probability at least $1 - d^{-1}$.

Proof of Lemma B.2. First observe that we have that \widehat{F} with I connected components has an edge set $E(\widehat{F}) = \{\widehat{E}_i\}_{i \in [d-I]}$, where $\{\widehat{E}_i\}_{i \in [d-1]}$ is an ordering of $E(\widehat{T})$ such that $|\widehat{\Theta}_{\widehat{E}_i}^{(1)}| \geq$

$|\widehat{\Theta}_{\widehat{E}_{i+1}}^{(1)}|$ for all $i \in [d-2]$. The remainder of the proof is similar to that of Lemma B.1 so we omit the details. \square

Proof of Corollary 3.3. Setting $r = 2 + \frac{C'C^4}{K}$, we get that

$$\theta = \max(rC'C^4/(r-2), rK)\sqrt{\log d/\lfloor n/2 \rfloor}.$$

Hence $\theta \geq rK\sqrt{\log d/\lfloor n/2 \rfloor}$, and a direct application of Lemma B.2 ensures that with high probability the edge set $E(\widehat{F})$ satisfies $\min_{e \in E(\widehat{F})} |\Theta_e^*| \geq (r-2)\theta/r$ when $\Theta^* \in \mathcal{S}_1(\theta, s)$ and $\min_{e \in E(\widehat{F})} |\Theta_e^*| = 0$ when $\Theta^* \in \mathcal{S}_0(s)$. The remaining part of the proposition follows directly by Proposition 3.1. \square

B.2 Cycle Detection Proofs

Lemma B.3. Assume the estimate $\widehat{\Theta}^{(1)}$ satisfies (3.3) on the parameter space $\mathcal{M}(s)$ and $\theta = rK\sqrt{\log d/\lfloor n/2 \rfloor}$ for some $r \geq 2$. Let \widehat{E} be the cycle obtained by Algorithm 4. We have $\min_{e \in \widehat{E}} |\Theta_e^*| \geq (r-2)\theta/r$ when $\Theta^* \in \mathcal{S}_1(\theta, s)$ and $\min_{e \in \widehat{E}} |\Theta_e^*| = 0$ when $\Theta^* \in \mathcal{S}_0(s)$ with probability at least $1 - 1/d$.

Proof of Lemma B.3. When $\Theta^* \in \mathcal{S}_1^1(\theta, s)$, there exists a cycle $E \subset E(G(\Theta^*))$. Hence by (B.1) when the algorithm terminates we will have covered only edges such that $|\widehat{\Theta}_e| > (r-1)\theta/r$. The latter, combined with (B.1) implies that until a cycle is reached, Algorithm 4 will have included only edges of true magnitude $|\Theta_e^*| \geq (r-2)\theta/r$, which completes the proof of the first claim.

On the other hand, when $\Theta^* \in \mathcal{S}_0(s)$, we have that there exists an edge $e \in \widehat{E}$ such that $\Theta_e^* = 0$, which completes the proof. \square

Proof of Corollary 3.4. To handle the null hypothesis observe that by Lemma B.3 we have $\min_{e \in \widehat{E}} |\Theta_e^*| = 0$. Invoking Proposition 3.1 gives the desired result. To show that the test is asymptotically powerful we use Proposition 3.1 in conjunction with the other result of Lemma B.3. \square

B.3 SAP Length Test Proofs

Lemma B.4. Assume the estimate $\widehat{\Theta}^{(1)}$ satisfies (3.3) on the parameter space $\mathcal{M}(s)$ and $\theta = rK\sqrt{\log d/\lfloor n/2 \rfloor}$ for some $r \geq 2$. If $\Theta^* \in \mathcal{S}_0(s)$, we have that $\min_{e \in \widehat{E}} |\Theta_e^*| = 0$. Conversely when $\Theta^* \in \mathcal{S}_1(\theta, s)$ we have $\min_{e \in \widehat{E}} |\Theta_e^*| \geq (r-2)\theta/r$ with probability at least $1 - 1/d$.

Proof of Lemma B.4. First, we consider the case when $\Theta^* \in \mathcal{S}_0(s)$. Since \widehat{E} is a SAP with $I+1$ edges and $G(\Theta^*)$ is free of such SAPs when $\Theta^* \in \mathcal{S}_0(s)$ we must have $\min_{e \in \widehat{E}} |\Theta_e^*| = 0$. Next, when $\Theta^* \in \mathcal{S}_1(\theta, s)$, by (B.1) $\min_{e \in \widehat{E}} |\widehat{\Theta}_e| > (r-1)\theta/r$. Thus by 3.3 $|\Theta_e^*| \geq (r-2)\theta/r$ with probability at least $1 - 1/d$. \square

Proof of Corollary 3.6. To handle the null hypothesis observe that by Lemma B.4 we are guaranteed to have $\min_{e \in \hat{E}} \Theta_e^* = 0$. Invoking Proposition 3.1 gives the desired result. To show that the test is asymptotically powerful we use Proposition 3.1 in conjunction with Lemma B.4. \square

B.4 Triangle-Free Tests Proofs

Lemma B.5. Assume the estimate $\hat{\Theta}^{(1)}$ satisfies (3.3) on the parameter space $\mathcal{M}(s)$ and $\theta = rK\sqrt{\log d/\lfloor n/2 \rfloor}$ for some $r \geq 2$. If $\Theta^* \in \mathcal{S}_0(s)$, we have that $\min_{e \in \hat{E}} |\Theta_e^*| = 0$. Conversely when $\Theta^* \in \mathcal{S}_1(\theta, s)$ we have $\min_{e \in \hat{E}} |\Theta_e^*| \geq (r-2)\theta/r$ with probability at least $1 - 1/d$.

Proof of Lemma B.5. Proof is the same as Lemma B.4, we omit the details. \square

Proof of Corollary 3.5. Proof is the same as Corollary 3.6, we omit the details. \square

B.5 Results for Thresholded Graphs

In this section, we discuss the hypothesis in (5.3)

$$\mathbf{H}_0 : G(\Theta^*, \mu) \text{ disconnected vs } \mathbf{H}_1 : G(\Theta^*, \mu) \text{ connected}$$

and prove the result on the test $\psi_{\alpha, \hat{T}}^B(\mathcal{D}_2)$ proposed in Section 5.3. Denote

$$\begin{aligned} \mathcal{S}_0(s) &:= \left\{ \Theta \in \mathcal{M}(s) : G(\Theta^*, \mu) \text{ disconnected} \right\} \text{ and,} \\ \mathcal{S}_1(\theta, s) &:= \left\{ \Theta \in \mathcal{M}(s) : G(\Theta^*, \mu + \theta) \text{ connected} \right\}. \end{aligned}$$

We have the following result parallel to Corollary 3.2.

Corollary B.6. Let $\theta = (2K + C'C^4)\log d/\lfloor n/2 \rfloor$ for an absolute constant $C' > 0$. Assume that $\mu \geq \theta$ and (3.10) holds. Then for any fixed level α , the test $\psi_{\alpha, \hat{E}}^B(\mathcal{D}_2)$ from the output of Algorithm 3 satisfies:

$$\limsup_{n \rightarrow \infty} \sup_{\Theta^* \in \mathcal{S}_0(s)} \mathbb{P}(\text{reject } \mathbf{H}_0) \leq \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\Theta^* \in \mathcal{S}_1(\theta, s)} \mathbb{P}(\text{reject } \mathbf{H}_0) = 1.$$

Proof. The proof is same as the one of Corollary 3.2. We only need to change the proof of Proposition 3.1 in Appendix E.1 to the thresholded version. In fact, we only need to change (E.10). When $|\Theta_{jk}^*| < \mu$, we have

$$\sqrt{n}|\tilde{\Theta}_{jk}| \geq \sqrt{n}|\kappa\sqrt{\log d/n}| - \sqrt{n} \max_{(j,k) \in E} |\tilde{\Theta}_{jk} - \Theta_{jk}^*| + \sqrt{n}\mu \geq \sqrt{n}\mu + c_{1-\alpha, E}.$$

All the remaining part of the argument is the same. \square

C Multi-Edge NAS Proofs

Before proceeding with the proofs we remind the reader definition (4.2) of $\mathbf{A}_{S,S'} = \mathbf{A}_0 + \mathbf{A}_S + \mathbf{A}_{S'}$.

Proof of Theorem 4.1. We denote by $\mathcal{L} := \exp(M(\mathbf{C}, d_{G_0}, \log |\mathbf{C}|))$ the cardinality of the $(\log |\mathbf{C}|)$ -packing. Using Proposition C.1 with the constants from Setting 1, and the fact that we have a $(\log |\mathbf{C}|)$ -packing it suffices to control:

$$\begin{aligned} & \frac{2}{\mathcal{L}^2} \sum_{(S,S'): d_{G_0}(S,S') \geq \log |\mathbf{C}|} \exp \left(\frac{n |\mathcal{V}_{S',S} \cap V(S)| \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 (2 \|\mathbf{A}_{S,S'}\|_2 \theta)^{2d_{G_0}(S,S')} \theta^2}{2d_{G_0}(S,S') + 2} \right) \\ & + \frac{2}{\mathcal{L}^2} \sum_{(S,S'): d_{G_0}(S,S') = 0} \exp \left(n |S \cap S'| \theta^2 + \frac{n |\mathcal{V}_{S',S} \cap V(S)| \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 (2 \|\mathbf{A}_{S,S'}\|_2 \theta)^2 \theta^2}{4} \right), \end{aligned}$$

When the cardinality of each $|S| \leq U$ the above expression can further be controlled by:

$$\underbrace{\frac{2}{\mathcal{L}^2} \sum_{(S,S'): d_{G_0}(S,S') \geq \log |\mathbf{C}|} \exp \left(\frac{2nU^3(2\Lambda_0\theta)^{2d_{G_0}(S,S')} \theta^2}{2d_{G_0}(S,S') + 2} \right)}_{I_1} + \underbrace{\frac{2}{\mathcal{L}} \exp \left(nU\theta^2 + 2nU^3(\Lambda_0\theta)^2 \theta^2 \right)}_{I_2},$$

where $\Lambda_0 := \|\mathbf{A}_0\|_2 + 2U$, and we used the facts that $|S \cap S'|, \|\mathbf{A}_S\|_2, \|\mathbf{A}_{S'}\|_2 \leq U$, $|\mathcal{V}_{S',S} \cap V(S)| \leq |V(S)| \leq 2U$ and that $d_{G_0}(S, S') = 0$ is only possible when $S \equiv S'$. We first deal with the term I_2 . For values of

$$\theta < \frac{1}{2U\Lambda_0},$$

we have $I_2 \leq \frac{2}{\mathcal{L}} \exp(2nU\theta^2)$. Hence when $\theta \leq \kappa \sqrt{\frac{\log \mathcal{L}}{nU}}$ for some sufficiently small $0 < \kappa$ we have $I_2 = o(1)$. Next, similarly to the proof of Step 4 in Theorem 2.1 (recall the definition of K_r) we can show:

$$I_1 \leq \left(1 - \frac{1}{\mathcal{L}}\right) + \frac{12(2\Lambda_0\theta)^{2\lfloor \log \mathcal{L} \rfloor + 2}}{1 - (2\Lambda_0\theta)^2} \mathcal{L}^{-1} \max_{\lfloor \log \mathcal{L} \rfloor + 1 \leq r} \frac{K_r}{r}.$$

Paying closer attention to the second term we have:

$$\begin{aligned} \frac{12(2\Lambda_0\theta)^{2\lfloor \log \mathcal{L} \rfloor + 2}}{1 - (2\Lambda_0\theta)^2} \mathcal{L}^{-1} \max_{\lfloor \log \mathcal{L} \rfloor + 1 \leq r} \frac{K_r}{r} & \leq 24(2\Lambda_0\theta)^{2\lfloor \log \mathcal{L} \rfloor + 2} \frac{\binom{\mathcal{L}}{2} + \mathcal{L}}{\mathcal{L}} \\ & \leq 24(2\Lambda_0\theta)^{2\lfloor \log \mathcal{L} \rfloor + 2} \mathcal{L} = o(1), \end{aligned}$$

with the last equalities hold for:

$$2\Lambda_0\theta \leq \frac{1}{2} < \exp(-1/2).$$

This completes the proof. \square

Proof of Theorem 4.2. Note that by the definition of \mathcal{B} we have:

$$\max_{S, S' \in \mathcal{C}} \left((\|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S, S'}\|_2^2) \wedge (|\mathcal{V}_{S, S'}| \|\mathbf{A}_{S, S'}\|_1^2) \right) \leq \mathcal{B}.$$

Therefore using Proposition C.1 it suffices to control:

$$\overline{D}_{\chi^2} := \frac{1}{|\mathcal{C}|^2} \sum_{S, S' \in \mathcal{C}} \exp \left[|\mathcal{V}_{S, S'} \cap V(S)| n \theta^2 (\mathcal{R} + \mathcal{B} \theta^2) \right].$$

First note that by $|\mathcal{V}_{S, S'} \cap V(S)| = \sum_{v \in V(S)} \mathbb{1}(v \in \mathcal{V}_{S, S'})$, we have:

$$\overline{D}_{\chi^2} \leq \frac{1}{|\mathcal{C}|^2} \sum_{S, S' \in \mathcal{C}} \exp \left[n \theta^2 (\mathcal{R} + \mathcal{B} \theta^2) \sum_{v \in V(S)} \mathbb{1}(v \in \mathcal{V}_{S, S'}) \right].$$

Denote by $\mathbb{P}_{S'}$ the measure induced by drawing S' uniformly from \mathcal{C} . Under the assumption: $\theta < \sqrt{\mathcal{R}/\mathcal{B}}$, and using the fact that the random variables $\{\mathbb{1}(v \in \mathcal{V}_{S, S'}) \mid v \in V(S)\}$ are negatively associated for every fixed $S \in \mathcal{C}$, we obtain:

$$\begin{aligned} \log \overline{D}_{\chi^2} &\leq \max_{S \in \mathcal{C}} \left[\sum_{v \in V(S)} \log \left[\exp(2\mathcal{R}n\theta^2) \mathbb{P}_{S'}(v \in \mathcal{V}_{S, S'}) + (1 - \mathbb{P}_{S'}(v \in \mathcal{V}_{S, S'})) \right] \right] \\ &\leq \max_{S \in \mathcal{C}} \left[(\exp(2\mathcal{R}n\theta^2) - 1) \sum_{v \in V(S)} \mathbb{P}_{S'}(v \in \mathcal{V}_{S, S'}) \right] \\ &\leq \exp(2\mathcal{R}n\theta^2) \max_{S \in \mathcal{C}} \mathbb{E}_{S'} |\mathcal{V}_{S, S'}|, \end{aligned}$$

where the expectation $\mathbb{E}_{S'}$ is taken with respect to a uniform draw of $S' \in \mathcal{C}$. The first inequality above is due to negative association, the second inequality is due to $\log(1+x) \leq x$. Hence for values of θ

$$\theta \leq \sqrt{\frac{\log \left([\max_{S \in \mathcal{C}} \mathbb{E}_{S'} |\mathcal{V}_{S, S'}|]^{1/2} [\max_{S \in \mathcal{C}} \mathbb{E}_{S'} |\mathcal{V}_{S, S'}|]^{-1} \right)}{2n\mathcal{R}}} = \sqrt{\frac{\log \left([\max_{S \in \mathcal{C}} \mathbb{E}_{S'} |\mathcal{V}_{S, S'}|]^{-1} \right)}{4n\mathcal{R}}},$$

we have $\overline{D}_{\chi^2} \leq \exp([\max_{S \in \mathcal{C}} \mathbb{E}_{S'} |\mathcal{V}_{S, S'}|]^{1/2})$, and therefore:

$$\liminf_n \gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) \geq 1 - \frac{1}{2} \sqrt{\exp([\max_{S \in \mathcal{C}} \mathbb{E}_{S'} |\mathcal{V}_{S, S'}|]^{1/2}) - 1} = 1,$$

where the last equality holds since $M_B(\mathcal{C}, \mathbf{V}) \rightarrow \infty$ implies $\max_{S \in \mathcal{C}} \mathbb{E}_{S'} |\mathcal{V}_{S, S'}| = o(1)$. \square

Proposition C.1. Let $G_0 \in \mathcal{G}_0$ and \mathcal{C} be a NAS set with null base G_0 . Then for any collection of vertex buffers $\mathbf{V} = \{\mathcal{V}_{S, S'}\}_{S, S' \in \mathcal{C}}$ and any of the following two choices:

Setting 1: $\mathcal{H}_{S, S'} := \frac{(|\mathcal{V}_{S, S'} \cap V(S)| \wedge |\mathcal{V}_{S, S'} \cap V(S')|) \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2}{\|\mathbf{A}_{S, S'}\|_2^2}$ and $\mathcal{K}_{S, S'} := 2\|\mathbf{A}_{S, S'}\|_2$;

Setting 2: $\mathcal{H}_{S, S'} := \frac{|\mathcal{V}_{S, S'} \cap V(S)| |\mathcal{V}_{S, S'} \cap V(S')|}{\|\mathbf{A}_{S, S'}\|_1^2}$ and $\mathcal{K}_{S, S'} := 2\|\mathbf{A}_{S, S'}\|_1$.

when the signal strength satisfies

$$\theta < \min_{S, S' \in \mathbf{C}} \frac{1 - C^{-1}}{2\sqrt{2}\|\mathbf{A}_{S, S'}\|_1}, \quad (\text{C.1})$$

we have:

$$\gamma(\mathcal{S}_0(\theta, s), \mathcal{S}_1(\theta, s)) \geq 1 - \frac{1}{2} \sqrt{\frac{1}{|\mathbf{C}|^2} \sum_{S, S' \in \mathbf{C}} \exp \left[n \left(|S \cap S'| \theta^2 + \frac{\mathcal{H}_{S, S'}(\mathcal{K}_{S, S'} \theta)^{2(d_{G_0}(S, S') \vee 1 + 1)}}{2(d_{G_0}(S, S') \vee 1 + 1)} \right) \right]} - 1,$$

Remark C.1. Proposition C.1 continues to hold for parameter classes $\mathcal{M}(s)$ not imposing bounds on the ℓ_1 norm of the precision matrix Θ , if we substitute $\|\mathbf{A}_{S, S'}\|_1$ with $\|\mathbf{A}_{S, S'}\|_2$ in (C.1). In fact tracking the proof, it is easy to see that the theorem also remains valid for parameter spaces such that $\mathbf{I} + \theta \mathbf{A}_0 \in \mathcal{S}_0(\theta, s)$ and $\mathbf{I} + \theta(\mathbf{A}_0 + \mathbf{A}_S) \in \mathcal{S}_1(\theta, s)$ for all $S \in \mathbf{C}$, with (C.1) replaced by $\theta < \min_{S, S' \in \mathbf{C}} \frac{1 - C^{-1}}{\|\mathbf{A}_{S, S'}\|_2 \vee \sqrt{2} \mathcal{K}_{S, S'}}$.

Definition C.1. Let \mathbf{C} be a NAS with null base $G_0 \in \mathcal{G}_0$ whose adjacency matrix is \mathbf{A}_0 . We call a set of constants $\{\mathcal{H}_{S, S'}, \mathcal{K}_{S, S'} : S, S' \in \mathbf{C}\}$ *admissible* with respect to the pair (G_0, \mathbf{C}) if for all even integers $k \geq 4$ the following holds:

$$\text{Tr}(\mathbf{A}_{S, S'}^k + \mathbf{A}_0^k - (\mathbf{A}_0 + \mathbf{A}_S)^k - (\mathbf{A}_0 + \mathbf{A}_{S'})^k) \leq \mathcal{H}_{S, S'} \mathcal{K}_{S, S'}^k, \quad (\text{C.2})$$

for all $S, S' \in \mathbf{C}$.

We will see that any admissible set of constants $\{\mathcal{H}_{S, S'}, \mathcal{K}_{S, S'} : S, S' \in \mathbf{C}\}$ yields the lower bound on the minimax risk claimed by Proposition C.1.

Proof of Proposition C.1. We will in fact show a slightly stronger result than presented, involving the following additional combinatorial quantity:

$$\Delta_{S, S'} := |\{\text{triangles in } G(\mathbf{A}_{S, S'}) \text{ with at least one edge in } S \text{ and another edge in } S'\}|.$$

Step 1 (Matrix Construction).

In this step we construct a set of precision matrices and argue that they fall into the parameter set $\mathcal{M}(s)$. Take $\Theta_0 = \mathbf{I} + \theta \mathbf{A}_0$, $\Theta_S = \mathbf{I} + \theta(\mathbf{A}_0 + \mathbf{A}_S)$, $\Theta_{S, S'} = \mathbf{I} + \theta \mathbf{A}_{S, S'}$, for $S, S' \in \mathbf{C}$ and some $\theta > 0$. For any $S, S' \in \mathbf{C}$ we have:

$$\begin{aligned} \max(\|\mathbf{A}_0\|_2, \|\mathbf{A}_0 + \mathbf{A}_S\|_2, \|\mathbf{A}_{S, S'}\|_2) &\leq \|\mathbf{A}_{S, S'}\|_2 \leq \|\mathbf{A}_{S, S'}\|_1 \\ \max(\|\mathbf{A}_0\|_1, \|\mathbf{A}_0 + \mathbf{A}_S\|_1, \|\mathbf{A}_{S, S'}\|_1) &\leq \|\mathbf{A}_{S, S'}\|_1, \end{aligned}$$

where these inequalities hold since all matrices \mathbf{A}_0 , $\mathbf{A}_0 + \mathbf{A}_S$ and $\mathbf{A}_{S, S'}$ consist only of non-negative entries.

Similarly to the proof of Theorem 2.1 we can make sure that the matrices Θ_0 and Θ_S fall into the set $\mathcal{M}(s)$ and in addition the matrix $\Theta_{S,S'}$ is strictly positive definite if $\theta < \frac{1-C^{-1}}{\Gamma}$. Thus by assumption the graphs $G(\Theta_0) \in \mathcal{G}_0$ and $G(\Theta_S) \in \mathcal{G}_1$ for all $S \in \mathcal{C}$. In addition this implies that the matrices $\Theta_0, \Theta_S, \Theta_{S,S'}$ are strictly positive definite for any $S, S' \in \mathcal{C}$.

Step 2 (*Risk and Trace Bounds*).

In this step we will lower bound the risk, and will further derive some combinatorial bounds on the traces of powers of adjacency matrices. These bounds are more detail tracking compared to bounds discussed in Theorem 2.1. Similarly to Step 2 of the proof of Theorem 2.1 it suffices to bound:

$$\begin{aligned} & \left(\frac{\det(\mathbf{I} + \theta(\mathbf{A}_0 + \mathbf{A}_S))}{\det(\mathbf{I} + \theta\mathbf{A}_0)} \right)^{n/2} \left(\frac{\det(\mathbf{I} + \theta(\mathbf{A}_0 + \mathbf{A}_{S'}))}{\det(\mathbf{I} + \theta(\mathbf{A}_{S,S'}))} \right)^{n/2} \\ &= \exp \left(\frac{n}{2} \sum_{k=1}^{\infty} \frac{(-\theta)^k}{k} \text{Tr} \left(\mathbf{A}_{S,S'}^k + \mathbf{A}_0^k - (\mathbf{A}_0 + \mathbf{A}_S)^k - (\mathbf{A}_0 + \mathbf{A}_{S'})^k \right) \right), \end{aligned}$$

Similarly to Step 3 of Theorem 2.1 it is easy to argue that for any $k \in \mathbb{N}$ we have:

$$\text{Tr}(\mathbf{A}_{S,S'}^k + \mathbf{A}_0^k - (\mathbf{A}_0 + \mathbf{A}_S)^k - (\mathbf{A}_0 + \mathbf{A}_{S'})^k) \geq 0.$$

We will consider three cases: (1) $k < 2(d_{G_0}(S, S') + 1)$, (2) $k < 4$ and $k \geq 2(d_{G_0}(S, S') + 1)$ and (3) $k \geq 4$ and $k \geq 2(d_{G_0}(S, S') + 1)$. For $k < 2(d_{G_0}(S, S') + 1)$, similarly to the argument in the Step 3 of the proof of Theorem 2.1, the above is in fact an equality. However, for the case $k < 4$ and $k \geq 2(d_{G_0}(S, S') + 1)$, we will instead show the following two more precise bounds:

$$\text{Tr}(\mathbf{A}_0^2 + \mathbf{A}_{S,S'}^2 - (\mathbf{A}_0 + \mathbf{A}_S)^2 - (\mathbf{A}_0 + \mathbf{A}_{S'})^2) \leq 4|S \cap S'|, \quad (\text{C.3})$$

$$\text{Tr}(\mathbf{A}_0^3 + \mathbf{A}_{S,S'}^3 - (\mathbf{A}_0 + \mathbf{A}_S)^3 - (\mathbf{A}_0 + \mathbf{A}_{S'})^3) \geq 6\Delta_{S,S'}. \quad (\text{C.4})$$

The left hand side of (C.3) contains edges lying only in the intersection $S \cap S'$ since all closed walks containing at least one edge in G_0 cancel out. By definition the number of such edges is $|S \cap S'|$. In addition, each closed walk of length 2 in (C.3), has precisely one edge in S and one edge in S' . Fixing the first edge to be in S and the second edge to be in S' we notice that each edge appears twice in the total count — once for each of its two vertices. Further multiplying by 2 to adjust for the ordering of the edges we obtain $4|S \cap S'|$. Next, notice that expression (C.4) contains only closed walks which are triangles. In addition, similarly to the logic above, only walks containing one edge from S and S' survive in (C.4). Further, each triangle is contained at least 6 times — once per each of its vertices and once per its 2 orientations. This completes the proof of (C.4).

We will check the third case in the next step by checking the admissibility in Definition C.1.

Step 3 (Verifying Admissibility).

In this step we show that both settings below are admissible for any pair (G_0, \mathbf{C}) :

$$\begin{aligned} \text{S1: } \mathcal{H}_{S,S'} &:= \frac{(|\mathcal{V}_{S,S'} \cap V(S)| \wedge |\mathcal{V}_{S,S'} \cap V(S')|) \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2}{\|\mathbf{A}_{S,S'}\|_2^2} \text{ and } \mathcal{K}_{S,S'} := 2\|\mathbf{A}_{S,S'}\|_2; \\ \text{S2: } \mathcal{H}_{S,S'} &:= \frac{|\mathcal{V}_{S,S'} \cap V(S)| |\mathcal{V}_{S,S'} \cap V(S')|}{\|\mathbf{A}_{S,S'}\|_1^2} \text{ and } \mathcal{K}_{S,S'} := 2\|\mathbf{A}_{S,S'}\|_1. \end{aligned}$$

Before we prove that the constants in Settings 1 and 2 are admissible, we will show a simple and general bound on closed walks over a sequence of graphs. Let $E_1, \dots, E_j \subset \bar{E}$ be fixed edge sets and $G_1 = (\bar{V}, E_1), G_2 = (\bar{V}, E_2), \dots, G_j = (\bar{V}, E_j)$ be graphs with vertex set \bar{V} , adjacency matrices $\mathbf{A}_1, \dots, \mathbf{A}_j$. Denote w_{ii} as the number of closed walks starting and ending at vertex i such that its ℓ^{th} edge locates on G_ℓ for $\ell \in [j]$. Note that w_{ii} , precisely equals to the (i, i) -th entry of the matrix \mathbf{A} , where $\mathbf{A} = \prod_{\ell \in [j]} \mathbf{A}_\ell$, i.e.

$$w_{ii} = A_{ii} = \left[\prod_{\ell \in [j]} \mathbf{A}_\ell \right]_{ii} \leq \|\mathbf{A}\|_2 \leq \prod_{\ell \in [j]} \|\mathbf{A}_\ell\|_2. \quad (\text{C.5})$$

We conclude that for any fixed vertex $i \in \bar{V}$, we have that the number of closed walks starting and ending at vertex i walking on the edges of G_ℓ for $\ell \in [j]$ is at most: $\prod_{\ell \in [j]} \|\mathbf{A}_\ell\|_2$.

Following the above argument, we will prove below the admissibility of the constants in Setting 1. More precisely we will prove the following

$$\begin{aligned} & \text{Tr}(\mathbf{A}_0^k + \mathbf{A}_{S,S'}^k - (\mathbf{A}_0 + \mathbf{A}_S)^k - (\mathbf{A}_0 + \mathbf{A}_{S'})^k) \\ & \leq \left[\frac{\frac{k^3}{4} (|\mathcal{V}_{S,S'} \cap V(S)| \wedge |\mathcal{V}_{S,S'} \cap V(S')|) \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2}{\|\mathbf{A}_{S,S'}\|_2^2} \right] \|\mathbf{A}_{S,S'}\|_2^k,^3 \end{aligned} \quad (\text{C.6})$$

which implies the admissibility of Setting 1, by the trivial bound $\frac{k^3}{4} \leq 2^k$. Recall that k is even.

To prove (C.6), we remind the reader that the trace of an adjacency matrix, counts the number of closed walks of length k in the graph. A walk will only be counted in the LHS of (C.6), if it contains an edge from the set S and another edge from the set S' . In the remainder of this step of the proof, we will bound the number of closed walks containing edges from both S and S' .

Denote $\mathcal{C}_{S,S'}^{(k)} = \{\text{closed walks } \mathcal{C} \text{ of length } k \text{ on } G(\mathbf{A}_{S,S'}), \text{ with edges } e, e' \in \mathcal{C}, e \neq e', e \in S, e' \in S'\}$. We will denote closed walks of length k by $\mathcal{C} = v_0^{\mathcal{C}} \rightarrow v_1^{\mathcal{C}} \rightarrow \dots \rightarrow v_{k-1}^{\mathcal{C}} \rightarrow v_0^{\mathcal{C}}$, where $v_j^{\mathcal{C}}$ is the j^{th} vertex of \mathcal{C} and $(v_j^{\mathcal{C}}, v_{j+1}^{\mathcal{C}})$ is its j^{th} edge, and indexation is taken modulo k . All indexation below will also be taken modulo k where applicable.

³Of note, one can show that in the special case of buffers $\mathcal{V}_{S,S'} = \{V(E_0 \cup S) \cap V(S')\} \cup \{V(E_0 \cup S') \cap V(S)\}$ a stronger bound holds with the constant $2\binom{k}{2}$ in place of $\frac{k^3}{4}$. This is however irrelevant for our purposes, since we relax this constant.




Figure 19: A closed walk \mathcal{C} from the set $\mathcal{C} \in \mathcal{C}_k(v, o, i, j)$.

For any $0 \leq j < k/2$, $0 \leq i \leq j$, and a vertex $v \in \bar{V}$, we first count the number of closed walks in the set $\mathcal{C}_k(v, o, i, j) = \{\mathcal{C} \in \mathcal{C}_{S, S'}^{(k)} \mid v_{i+o}^{\mathcal{C}} = v, (v_{o-1}^{\mathcal{C}}, v_o^{\mathcal{C}}) \in S \text{ and } (v_{o+j}^{\mathcal{C}}, v_{o+j+1}^{\mathcal{C}}) \in S'\}$. Following (C.5), we have

$$|\mathcal{C}_k(v, o, i, j)| \leq \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S, S'}\|_2^{k-2}. \quad (\text{C.7})$$

Similarly, for $0 \leq j < k/2$, $0 \leq i \leq j$, and a vertex $v \in \bar{V}$ we define $\mathcal{C}'_k(v, i, j) = \{\mathcal{C} \in \mathcal{C}_{S, S'}^{(k)} \mid v_i^{\mathcal{C}} = v, (v_{o-1}^{\mathcal{C}}, v_o^{\mathcal{C}}) \in S' \text{ and } (v_{o+j}^{\mathcal{C}}, v_{o+j+1}^{\mathcal{C}}) \in S\}$. We also have

$$|\mathcal{C}'_k(v, o, i, j)| \leq \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S, S'}\|_2^{k-2}.$$

Define the sets $\mathcal{C}_k(o, j) = \{\mathcal{C} \in \mathcal{C}_{S, S'}^{(k)} \mid (v_{o-1}^{\mathcal{C}}, v_o^{\mathcal{C}}) \in S \text{ and } (v_{o+j}^{\mathcal{C}}, v_{o+j+1}^{\mathcal{C}}) \in S'\}$ and $\mathcal{C}'_k(j) = \{\mathcal{C} \in \mathcal{C}_{S, S'}^{(k)} \mid (v_{o-1}^{\mathcal{C}}, v_o^{\mathcal{C}}) \in S' \text{ and } (v_{o+j}^{\mathcal{C}}, v_{o+j+1}^{\mathcal{C}}) \in S\}$, so that we clearly have $\mathcal{C}_k(v, o, i, j) \subseteq \mathcal{C}_k(o, j)$ and $\mathcal{C}'_k(v, o, i, j) \subseteq \mathcal{C}'_k(o, j)$.

By the definition of the vertex buffer set $\mathcal{V}_{S, S'}$, each closed walk \mathcal{C} in $\mathcal{C}_k(v, o, i, j)$, has at least one vertex belonging to $\mathcal{V}_{S, S'|S}$ (and one vertex in $\mathcal{V}_{S, S'|S'}$). What is more, \mathcal{C} has a vertex in $\mathcal{V}_{S, S'|S}$ which belongs to the set $\{v_o, v_{o+1}, \dots, v_{o+j}\}$, since $v_o \in V(S)$ and $v_{o+j} \in V(S')$. This argument shows that:

$$\mathcal{C}_k(o, j) = \bigcup_{v \in \bar{V}} \bigcup_{0 \leq i \leq j} \mathcal{C}_k(v, o, i, j) \subseteq \bigcup_{v \in \mathcal{V}_{S, S'|S}} \bigcup_{0 \leq i \leq j} \mathcal{C}_k(v, o, i, j).$$

Hence:

$$|\mathcal{C}_k(o, j)| \leq (j+1) |\mathcal{V}_{S, S'|S}| \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S, S'}\|_2^{k-2},$$

and therefore by symmetry:

$$|\mathcal{C}_k(o, j)| \leq (j+1) (|\mathcal{V}_{S, S'|S}| \wedge |\mathcal{V}_{S, S'|S'}|) \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S, S'}\|_2^{k-2}, \quad (\text{C.8})$$

Similarly we have

$$|\mathcal{C}'_k(o, j)| \leq (j+1) (|\mathcal{V}_{S, S'|S}| \wedge |\mathcal{V}_{S, S'|S'}|) \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S, S'}\|_2^{k-2}. \quad (\text{C.9})$$

We now observe that:

$$\mathcal{C}_{S,S'}^{(k)} \subseteq \bigcup_{o=0}^{k-1} \bigcup_{j=0}^{k/2-1} \mathcal{C}_k(o, j) \cup \mathcal{C}'_k(o, j)$$

Notice that in the special case when $j = k/2 - 1$, we have $\bigcup_{o=0}^{k-1} \mathcal{C}_k(o, k/2 - 1) = \bigcup_{o=0}^{k-1} \mathcal{C}'_k(o, k/2 - 1)$. Indeed, by the fact that indexation is taken modulo k we trivially have $\mathcal{C}_k(o, k/2 - 1) = \mathcal{C}'_k(k/2 - o, k/2 - 1)$. Hence:

$$\mathcal{C}_{S,S'}^{(k)} \subseteq \bigcup_{o=0}^{k-1} \left[\bigcup_{j=0}^{k/2-2} \mathcal{C}_k(o, j) \cup \mathcal{C}_k(o, j) \right] \cup \mathcal{C}_k(o, k/2 - 1).$$

Applying (C.8) and (C.9), and summing over j we have:

$$\begin{aligned} |\mathcal{C}_{S,S'}^{(k)}| &\leq (|\mathcal{V}_{S,S'|S}| \wedge |\mathcal{V}_{S,S'|S'}|) \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S,S'}\|_2^{k-2} k \left[2 \sum_{j=1}^{\frac{k}{2}-1} j + \frac{k}{2} \right] \\ &= \frac{k^3}{4} (|\mathcal{V}_{S,S'|S}| \wedge |\mathcal{V}_{S,S'|S'}|) \|\mathbf{A}_S\|_2 \|\mathbf{A}_{S'}\|_2 \|\mathbf{A}_{S,S'}\|_2^{k-2}, \end{aligned}$$

which completes the proof of (C.6).

To check the admissibility for Setting 2, just as before we must have two edges in S and S' respectively within a closed walk of length k which is counted in the LHS of (C.2). Then, by the definition of vertex buffer set $\mathcal{V}_{S,S'}$, we certainly have two vertices in the set $\mathcal{V}_{S,S'}$ (one in $\mathcal{V}_{S,S'} \cap V(S)$ and one in $\mathcal{V}_{S,S'} \cap V(S')$). Notice that each vertex on the path is of degree at most $\|\mathbf{A}_{S,S'}\|_1$, and hence can give rise to at most $\|\mathbf{A}_{S,S'}\|_1$ continuations of the path. Therefore, having fixed two vertices from the sets $\mathcal{V}_{S,S'} \cap V(S)$ and $\mathcal{V}_{S,S'} \cap V(S')$, we are left with at most $\|\mathbf{A}_{S,S'}\|_1^{k-2}$ paths. Taking into account that we can position the two vertices on at most $k(k-1) \leq 2^k$ spots completes the proof.

Step 4 (Proof Completion).

In this step we complete the proof by arguing that if we are given any set of admissible constants $\{\mathcal{H}_{S,S'}, \mathcal{K}_{S,S'}\}_{S,S' \in \mathcal{C}}$ the bound on the minimax risk from the Theorem statement follows.

$$\begin{aligned} &\sum_{k=1}^{\infty} \theta^k \text{Tr} \left(\mathbf{A}_{S,S'}^k + \mathbf{A}_0^k - (\mathbf{A}_0 + \mathbf{A}_S)^k - (\mathbf{A}_0 + \mathbf{A}_{S'})^k \right) / k \\ &\leq 2|S \cap S'| \theta^2 - 2\Delta_{S,S'} \theta^3 + \mathcal{H}_{S,S'} \sum_{2|k, k \geq 2(d_{G_0}(S,S') \vee 1 + 1)} \frac{(\mathcal{K}_{S,S'} \theta)^k}{k} \\ &\leq 2|S \cap S'| \theta^2 - 2\Delta_{S,S'} \theta^3 + \frac{2\mathcal{H}_{S,S'} (\mathcal{K}_{S,S'} \theta)^{2(d_{G_0}(S,S') \vee 1 + 1)}}{2(d_{G_0}(S,S') \vee 1 + 1)}, \end{aligned}$$

where in the last inequality we used that $(\mathcal{K}_{S,S'}\theta)^2 \leq \frac{1}{2}$ by (C.1). In step 3 we verified that the constants $\{\mathcal{H}_{S,S'}, \mathcal{K}_{S,S'}\}$ given in the Theorem statement are admissible for (G_0, \mathbf{C}) in the sense of Definition C.1. This completes the proof. \square

C.1 Star Graph/Maximum Degree Proofs

Lemma C.2. Assume the estimate $\widehat{\Theta}^{(1)}$ satisfies (3.3) on the parameter space $\mathcal{M}(s)$ and $\theta = rK\sqrt{\log d/\lfloor n/2 \rfloor}$ for some $r \geq 2$. If $\Theta^* \in \mathcal{S}_0(s)$, we have that $\min_{e \in \widehat{E}} |\Theta_e^*| = 0$. Conversely when $\Theta^* \in \mathcal{S}_1(\theta, s)$ we have $\min_{e \in \widehat{E}} |\Theta_e^*| \geq (r-2)\theta/r$ with probability at least $1 - 1/d$.

Proof of Lemma C.2. First, we consider the case when $\Theta^* \in \mathcal{S}_0(s)$. Since \widehat{E} is a set of s edges we must have $\min_{e \in \widehat{E}} |\Theta_e^*| = 0$. Next, when $\Theta^* \in \mathcal{S}_1(\theta, s)$, by (B.1) $\min_{e \in \widehat{E}} |\widehat{\Theta}_e| > (r-1)\theta/r$. Thus by assumption (3.3) $|\Theta_e^*| \geq (r-2)\theta/r$ with probability at least $1 - 1/d$. \square

Proof of Corollary 4.3. To handle the null hypothesis observe that by Lemma C.2 we are guaranteed to have $\min_{e \in \widehat{E}} |\Theta_e^*| = 0$. Invoking Proposition 3.1 gives the desired result. To show that the test is asymptotically powerful we use Proposition 3.1 in conjunction with Lemma C.2. \square

C.2 Clique Detection Upper Bound Proof

Lemma C.3. Under \mathbf{H}_0 we have:

$$\sqrt{\widehat{\lambda}_{\min}} \geq 1 - \sqrt{\frac{s}{n}} - \sqrt{2 \frac{s \log(ed/s) + \log \alpha^{-1}}{n}} \geq 1 - (\sqrt{2} + 1) \sqrt{\frac{s \log(ed/s) + \log \alpha^{-1}}{n}}, \quad (\text{C.10})$$

with probability no less than $1 - 2\alpha$.

Proof of Lemma C.3. The proof relies on an implication of Gordon's comparison theorem for Gaussian processes, see (Vershynin, 2012, Corollary 5.35 e.g.). By this result we have that:

$$\sqrt{\lambda_{\min}(\widehat{\Sigma}_{CC})} \geq 1 - \sqrt{\frac{s}{n}} - t,$$

with probability at least $1 - 2\exp(-nt^2/2)$ for any $t \geq 0$. Using the union bound in conjunction to the standard inequality $\binom{d}{s} \leq \left(\frac{ed}{s}\right)^s$, by setting $t = \sqrt{2 \frac{s \log(ed/s) + \log \alpha^{-1}}{n}}$, we can ensure that (C.10) holds. \square

Lemma C.4. Under \mathbf{H}_1 we have:

$$\sqrt{\widehat{\lambda}_{\min}} \leq \frac{1 + \sqrt{\frac{s}{n}} + \sqrt{\frac{2 \log \alpha^{-1}}{n}}}{\sqrt{1 + (s-1)\theta}} \leq \frac{1 + (\sqrt{2} + 1) \sqrt{\frac{s \log(ed/s) + \log \alpha^{-1}}{n}}}{\sqrt{1 + (s-1)\theta}}, \quad (\text{C.11})$$

with probability at least $1 - 2\alpha$.

Proof of Lemma C.4. Taking in mind that $\Sigma = \Theta^{-1}$ by a simple calculation one can verify that for the set $C^* = \text{supp}(\mathbf{v})$, $\lambda_{\min}(\Sigma_{C^*C^*}) = [\lambda_{\max}(\Theta_{C^*C^*})]^{-1} = \frac{1}{1+(s-1)\theta}$ with a corresponding eigenvector $\frac{\mathbf{v}_{C^*}}{\sqrt{s}}$. Again using Corollary 5.35 of Vershynin (2012), and the fact that for two symmetric psd matrices \mathbf{A}, \mathbf{B} we have $\lambda_{\min}(\mathbf{A}\mathbf{B}\mathbf{A}) \leq \lambda_{\min}(\mathbf{A})^2 \lambda_{\max}(\mathbf{B})$, we have:

$$\sqrt{\lambda_d(\widehat{\Sigma}_{C^*C^*})} \leq \frac{1 + \sqrt{\frac{s}{n}} + \sqrt{\frac{2 \log \alpha^{-1}}{n}}}{\sqrt{1 + (s-1)\theta}},$$

with probability at least $1 - 2\alpha$. □

Proof of Proposition 4.4. Combining the results of Lemma C.3 and setting $\alpha = d^{-1}$ in Lemma C.4 it suffices to show there will be a gap between the bounds in (C.10) and (C.11). By simple algebra when

$$\theta > \kappa \sqrt{\frac{\log(ed/s)}{sn}},$$

for a sufficiently large κ the gap between (C.10) and (C.11) is implied, which completes the proof. □

D Auxiliary Results

Recall that for a symmetric matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ we denote its eigenvalues in decreasing with $\lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \dots \geq \lambda_d(\mathbf{A})$.

Theorem D.1 (Lidskii's Inequality (Helmke and Rosenthal, 1995)). For two symmetric $m \times m$ matrices \mathbf{A} and \mathbf{B} we have:

$$\sum_{j=1}^k \lambda_{i_j}(\mathbf{A} + \mathbf{B}) \leq \sum_{j=1}^k \lambda_{i_j}(\mathbf{A}) + \sum_{j=1}^k \lambda_j(\mathbf{B}),$$

for any $1 \leq i_1 < \dots < i_k \leq m$.

We now derive the following Corollary of Theorem D.1, which can also be viewed as a more general formulation of the latter result.

Corollary D.2. Under the setting of Theorem D.1, for any constants $c_1 \geq c_2 \geq \dots \geq c_m$, and a permutation σ on $\{1, \dots, m\}$ we have:

$$\sum_{j=1}^m c_{\sigma(j)} \lambda_j(\mathbf{A} + \mathbf{B}) \leq \sum_{j=1}^m c_{\sigma(j)} \lambda_j(\mathbf{A}) + \sum_{j=1}^m c_j \lambda_j(\mathbf{B}),$$

Proof of Corollary D.2. First notice that since

$$\sum_{j=1}^m \lambda_j(\mathbf{A} + \mathbf{B}) = \sum_{j=1}^m \lambda_j(\mathbf{A}) + \sum_{j=1}^m \lambda_j(\mathbf{B}),$$

it suffices to show the bound when for all j we have $c_j \geq 0$. Let $C = \{j : c_j \neq 0\}$. We will induct on the cardinality $|C|$, which we denote with k . When $k = 1$, the inequality immediately follows by Theorem D.1. Suppose the inequality holds for $k = l$. Notice that the inequality can equivalently be expressed as

$$\sum_{j=1}^m c_j \lambda_{\sigma^{-1}(j)}(\mathbf{A} + \mathbf{B}) \leq \sum_{j=1}^m c_j \lambda_{\sigma^{-1}(j)}(\mathbf{A}) + \sum_{j=1}^m c_j \lambda_j(\mathbf{B}), \quad (\text{D.1})$$

To see why it holds for $k = l + 1$ let $c^* = \min_{j \in C} c_j$, and notice that by Theorem D.1 we have:

$$\sum_{j \in C} c^* \lambda_{\sigma^{-1}(j)}(\mathbf{A} + \mathbf{B}) \leq \sum_{j \in C} c^* \lambda_{\sigma^{-1}(j)}(\mathbf{A}) + \sum_{j \in C} c^* \lambda_j(\mathbf{B}).$$

This implies that inequality (D.1) follows by an inequality in which we subtract c^* from all $c_j, j \in C$. This completes the proof by the induction hypothesis. \square

E Supplementary Material

For a real valued random variable X and a random $\ell \geq 1$, we define the Orlicz-norm ψ_ℓ as:

$$\|X\|_{\psi_\ell} = \sup_{p \geq 1} p^{-1/\ell} (\mathbb{E}|X|^p)^{1/p}. \quad (\text{E.1})$$

We mainly use the ψ_1 and ψ_2 norms. Recall that random variables with bounded ψ_1 and ψ_2 norms are called *sub-exponential* and *sub-Gaussian* (Vershynin, 2012, e.g.).

E.1 Minimal Structure Certification

Lemma E.1. The random variable $\mathbf{X} \sim N(0, (\boldsymbol{\Theta}^*)^{-1})$ satisfies:

$$\sup_{\|\alpha\|_2=1} \|\alpha^T \mathbf{X}\|_{\psi_2} \leq \|Z\|_{\psi_2} C < C, \quad (\text{E.2})$$

where $Z \sim N(0, 1)$.

Proof of Lemma E.1. The first part follows since we are assuming $\|\boldsymbol{\Theta}^{*-1}\|_2 \leq C$, and the second inequality follows via a direct calculation. \square

Lemma E.2. We have

$$\inf_{j,k \in [d]} \Delta_{jk} \geq C^{-2}.$$

Proof of Lemma E.2. By Issrelis' theorem we have $\Delta_{jk} = \Theta_{jj}^* \Theta_{kk}^* + \Theta_{jk}^* \Theta_{kj}^* \geq \Theta_{jj}^* \Theta_{kk}^* \geq C^{-2} > 0$. \square

Lemma E.3 (Covariance Concentration). There exists a universal constant $R > 0$, such that:

$$\|\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}^*\|_{\max} \leq RC^2 \sqrt{\log d/n},$$

with probability at least $1 - 1/d$ uniformly over the parameter space \mathcal{M} .

Proof of Lemma E.3. The proof follows by applying the inequality

$$\|X_i X_j\|_{\psi_1} \leq 2 \|X_i\|_{\psi_2} \|X_j\|_{\psi_2} \leq 2c^{-2}$$

(see Lemma E.1) in combination with Proposition 5.16 (Vershynin, 2012) and the union bound. \square

Lemma E.4. For large enough n and d we have:

$$\sup_{\boldsymbol{\Theta}^* \in \mathcal{M}(s)} \mathbb{P}(\max_{j,k \in [d]} \sqrt{n} |(\tilde{\Theta}_{jk} - \Theta_{jk}^*) + \boldsymbol{\Theta}_{*j}^{*T} (\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}^*) \boldsymbol{\Theta}_{*k}^*| > \xi_1) < \xi_2,$$

where $\xi_1 = \Xi_1 s \log d / \sqrt{n}$, $\xi_2 = 2/d$, where Ξ_1 is an absolute constant depending solely on K, L and C .

Proof of Lemma E.4. By elementary algebra we obtain the following representation:

$$\begin{aligned}
(\tilde{\Theta}_{jk} - \Theta_{jk}^*) + \Theta_{*j}^{*T}(\hat{\Sigma} - \Sigma^*)\Theta_{*k}^* &= - \underbrace{\frac{\hat{\Theta}_{*j}^T \hat{\Sigma}_{*\setminus j}(\hat{\Theta}_{\setminus j,k} - \Theta_{\setminus j,k}^*) + (\hat{\Theta}_{*j} - \Theta_{*j}^*)^T(\hat{\Sigma} - \Sigma^*)\Theta_{*k}^*}{\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j}}}_{I_1^{jk}} \\
&\quad - \underbrace{\Theta_{*j}^{*T}(\hat{\Sigma} - \Sigma^*)\Theta_{*k}^* \left[\frac{1}{\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j}} - 1 \right]}_{I_2^{jk}}, \tag{E.3}
\end{aligned}$$

where indexing with $\setminus j$ means dropping the corresponding column or element from the matrix or vector respectively. We deal with the terms I_1^{jk} first. By the triangle inequality followed by Hölder's inequality we have:

$$\max_{j,k \in [d]} |I_1^{jk}| \leq \max_{j,k \in [d]} \{ |\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j}|^{-1} (\|\hat{\Theta}_{*j}^T \hat{\Sigma}_{*\setminus j}\|_\infty \|\hat{\Theta}_{\setminus j,k} - \Theta_{\setminus j,k}^*\|_1 + \|\hat{\Theta}_{*j} - \Theta_{*j}^*\|_1 \|\hat{\Sigma} - \Sigma^*\|_{\max} \|\Theta_{*k}^*\|_1) \}$$

Since it is assumed that $\log d/n = o(1)$, take d and n large enough so that $\log d/n \leq 1/(4K^2)$. Let E be the event where (3.3) and (3.4) hold. Then on E we have:

$$\max_{j \in [d]} |\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j} - 1| \leq K \sqrt{\log d/n}.$$

This implies that on the same event $\sup_{j \in [d]} |\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j}|^{-1} \leq 2$. Continuing our bounds on the event E we have:

$$\|\hat{\Theta}_{*j}^T \hat{\Sigma}_{*\setminus j}\|_\infty \|\hat{\Theta}_{\setminus j,k} - \Theta_{\setminus j,k}^*\|_1 \leq K^2 s \log d/n = o(n^{-1/2}).$$

Furthermore on the intersection of E and the event from Lemma E.3, which holds with probability no less than $1 - 2/d$, we have:

$$\|\hat{\Theta}_{*j} - \Theta_{*j}^*\|_1 \|\hat{\Sigma} - \Sigma^*\|_{\max} \|\Theta_{*k}^*\|_1 \leq RC^2 K L s \log d/n = o(n^{-1/2}).$$

Combining the last four bounds we conclude that $\max_{j,k \in [d]} |I_1^{jk}| = o_p(n^{-1/2})$. Next we handle the terms $|I_2^{jk}|$. We have:

$$\max_{j,k \in [d]} |I_2^{jk}| \leq \max_{j,k \in [d]} \left\{ |\Theta_{*j}^{*T}(\hat{\Sigma} - \Sigma^*)\Theta_{*k}^*| \left| \frac{1}{\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j}} - 1 \right| \right\}.$$

Clearly on the event E under the assumption $\log d/n \leq 1/(4K^2)$, we have that:

$$\max_{j \in [d]} \left| \frac{1}{\hat{\Theta}_{*j}^T \hat{\Sigma}_{*j}} - 1 \right| \leq 2K \sqrt{\log d/n}. \tag{E.4}$$

Next we consider the random variables $\Theta_{*j}^{*T} \mathbf{X}^{\otimes 2} \Theta_{*k}^*$ for all $j, k \in [d]$. By Lemma E.1 we have:

$$\|\Theta_{*j}^{*T} \mathbf{X}^{\otimes 2} \Theta_{*k}^*\|_{\psi_1} \leq 2\|\Theta_{*j}^{*T} \mathbf{X}\|_{\psi_2} \|\mathbf{X}^T \Theta_{*k}^*\|_{\psi_2} \leq 2 \sup_{j \in [d]} \|\Theta_{*j}^*\|_2^2 C^2 \leq 2C^4,$$

with the last inequality following by the fact that $\sup_{j \in [d]} \|\Theta_{*j}^*\|_2^2 \leq \|\Theta^*\|_2^2 \leq C^2$ since $\Theta^* \in \mathcal{M}(s)$. Clearly then $\|\Theta_{*j}^{*T} \mathbf{X}^{\otimes 2} \Theta_{*k}^* - \Theta_{jk}^*\|_{\psi_1} \leq 4C^4$. Finally by the union bound and Proposition 5.16 (Vershynin, 2012), one concludes that there exists an absolute constant \tilde{C} , such that when $\sqrt{\log d/n} \leq 4C^4$, we have:

$$\max_{j,k \in [d]} |\Theta_{*j}^{*T} (\hat{\Sigma} - \Sigma^*) \Theta_{*k}^*| \leq 4\tilde{C}C^4 \sqrt{\log d/n}, \quad (\text{E.5})$$

with probability at least $1 - 2/d$. Combining this bound with (E.4) and our conditions we conclude that $\max_{j,k \in [d]} |I_2^{jk}| = o_p(n^{-1/2})$. This completes the proof. \square

Lemma E.5. Let $\{\mathbf{X}_i\}_{i=1}^n$ be identical (not necessarily independent), d -dimensional sub-Gaussian vectors with $\max_{i \in [n], j \in [d]} \|X_{ij}\|_{\psi_2} = C$. Then there exists an absolute constant $U > 0$ depending solely on C such that :

$$\max_{i \in [n]} \|\mathbf{X}_i^{\otimes 2}\|_{\max} < U \log(nd),$$

with probability at least $1 - 1/d$.

Proof of Lemma E.5. The proof is trivial upon noting that $\max_{i \in [n]} \|\mathbf{X}_i^{\otimes 2}\|_{\max} = \max_{i \in [n]} \|\mathbf{X}_i\|_{\infty}^2$ and combining (5.10) in Vershynin (2012) with the union bound. We omit the details. \square

Lemma E.6. Assuming that $s\sqrt{\log d/n}\sqrt{\log(nd)} = o(1)$, we have that for any fixed $j, k \in [d]$:

$$\lim_n \inf_{\Theta^* \in \mathcal{M}(s)} \mathbb{P}(|\hat{\Delta}_{jk} - \Delta_{jk}| \leq \tau(n)) = 1,$$

where $\tau(n) = o(1)$.

Proof of Lemma E.6. We start by showing that:

$$I := \max_{j,k \in [d]} n^{-1} \sum_{i=1}^n (\hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*k} - \hat{\Theta}_{jk} - \Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* + \Theta_{jk}^*)^2,$$

is asymptotically negligible. Since $(a - b)^2 \leq 2(a^2 + b^2)$, we have that:

$$|I| \leq \underbrace{2 \max_{j,k \in [d]} n^{-1} \left[\sum_{i=1}^n (\hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*k} - \Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^*)^2 \right]}_{I_1} + \underbrace{2 \max_{j,k \in [d]} (\hat{\Theta}_{jk} - \Theta_{jk}^*)^2}_{I_2}.$$

Clearly by assumption (3.3)

$$|I_2| \leq K^2 \log d/n,$$

with probability at least $1 - 1/d$. Next we handle I_1 . By the triangle inequality the following holds:

$$I_1^{1/2} \leq \underbrace{\max_{j,k \in [d]} \left[n^{-1} \sum_{i=1}^n (\hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} (\hat{\Theta}_{*k} - \Theta_{*k}^*))^2 \right]^{1/2}}_{I_{11}} + \underbrace{\max_{j,k \in [d]} \left[n^{-1} \sum_{i=1}^n ((\hat{\Theta}_{*j} - \Theta_{*j}^*)^T \mathbf{X}_i^{\otimes 2} \Theta_{*k}^*)^2 \right]^{1/2}}_{I_{12}}.$$

We first handle I_{11}^2 :

$$I_{11}^2 = \max_{j,k \in [d]} (\hat{\Theta}_{*k} - \Theta_{*k}^*)^T \underbrace{\frac{1}{n} \sum_{i=1}^n \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*j}^{\otimes 2} \mathbf{X}_i^{\otimes 2}}_{\mathbf{M}} (\hat{\Theta}_{*k} - \Theta_{*k}^*).$$

Using Lemma E.5, we can handle \mathbf{M} in the following way:

$$\|\mathbf{M}\|_{\max} \leq \max_{i \in [d]} \|\mathbf{X}_i^{\otimes 2}\|_{\max} \max_{j \in [d]} \frac{1}{n} \sum_{i=1}^n \hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*j} \leq O_p(\log(nd)) \max_{j \in [d]} \|\hat{\Theta}_{*j}\|_1 \|\hat{\Sigma} \hat{\Theta}_{*j}\|_{\infty}. \quad (\text{E.6})$$

By the definition of $\hat{\Theta}_{*j}$ we have: $\max_{j \in [d]} \|\hat{\Sigma} \hat{\Theta}_{*j}\|_{\infty} \leq (1 + K \sqrt{\log d/n})$. Hence:

$$\|\mathbf{M}\|_{\max} \leq O_p(\log(nd)) \max_{j \in [d]} (\|\Theta_{*j}^*\|_1 + \|\hat{\Theta}_{*j} - \Theta_{*j}^*\|_1) (1 + K \sqrt{\log d/n}) = O_p(\log(nd)) L,$$

where we used assumption (3.4). Thus:

$$|I_{11}| \leq \max_{k \in [d]} \|\hat{\Theta}_{*k} - \Theta_{*k}^*\|_1^2 L O_p(\log(nd)) = O_p(s^2 \log d/n \log(nd)) = o_p(1),$$

with probability at least $1 - 2/d$ uniformly over $\mathcal{M}(s)$. By a similar argument we can show that I_{12} is of the same order. Putting everything together we conclude:

$$I = O_p(s^2 \log d \log(nd)/n) = o_p(1). \quad (\text{E.7})$$

Next, we argue that for any fixed j, k we have that the difference $n^{-1} \sum_{i=1}^n (\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*j}^* - \Theta_{*k}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^*)^2 - \mathbb{E}(\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{*k}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*j}^*)^2$ is small. Since as in Lemma E.4 we have $\|\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^*\|_{\psi_1} \leq 2C^4$, then certainly $\text{Var}((\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{*k}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*j}^*)^2)$ is finite for any fixed $j, k \in [d]$. A usage of Chebyshev's inequality gives us that:

$$n^{-1} \sum_{i=1}^n (\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*j}^* - \Theta_{*k}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^*)^2 - \mathbb{E}(\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{*k}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*j}^*)^2 = o_p(1).$$

Finally note that by the triangle inequality the following two inequalities hold for any fixed $j, k \in [d]$:

$$\begin{aligned} \left[n^{-1} \sum_{i=1}^n (\hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*k} - \hat{\Theta}_{jk})^2 \right]^{1/2} &\leq \left[n^{-1} \sum_{i=1}^n (\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{jk}^*)^2 \right]^{1/2} + I^{1/2}, \\ \left[n^{-1} \sum_{i=1}^n (\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{jk}^*)^2 \right]^{1/2} &\leq \left[n^{-1} \sum_{i=1}^n (\hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*k} - \hat{\Theta}_{jk})^2 \right]^{1/2} + I^{1/2}. \end{aligned}$$

Observe that $n^{-1} \sum_{i=1}^n (\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{jk}^*)^2 = \mathbb{E}(\Theta_{*j}^{*T} \mathbf{X}^{\otimes 2} \Theta_{*k}^* - \Theta_{jk}^*)^2 + o_p(1) = O_p(1)$. This completes the proof. \square

Lemma E.7. Let $X_n(\mathbf{r})$ and $\xi_n(\mathbf{r})$ be two sequences of random variables, depending on a parameter $\mathbf{r} \in \mathbf{R}$. Suppose that $\lim_n \sup_{\mathbf{r} \in \mathbf{R}} \sup_t |\mathbb{P}(X_n(\mathbf{r}) \leq t) - F(t)| = 0$ where F is a continuous cdf, and $\lim_n \inf_{\mathbf{r} \in \mathbf{R}} \mathbb{P}(|1 - \xi_n(\mathbf{r})| \leq \tau(n)) = 1$ for $\tau(n) = o(1)$. Assume in addition that F is Lipschitz, i.e., there exist $\kappa > 0$, such that for any $t, s \in \mathbb{R} : |F(t) - F(s)| \leq \kappa|t - s|$. Then we have:

$$\lim_n \sup_{\mathbf{r} \in \mathbf{R}} \sup_t \left| \mathbb{P}\left(\frac{X_n(\mathbf{r})}{\xi_n(\mathbf{r})} \leq t\right) - F(t) \right| = 0.$$

Proof. The proof goes through by a direct argument. We omit the details. \square

Proof of Proposition 3.1. The proof of the strong error rate control is based upon verifying conditions from Theorem 5.1 of Chernozhukov et al. (2013). The rate $(\log(dn))^6/n = o(1)$ in (3.10) holds due to recent advances in a Berry-Esseen bound due to Chernozhukov et al. (2016). Define:

$$\begin{aligned} \Upsilon_1 &:= \sqrt{n} \max_{(j,k) \in E} |(\tilde{\Theta}_{jk} - \Theta_{jk}^*) + \Theta_{*j}^{*T}(\hat{\Sigma} - \Sigma^*)\Theta_{*k}^*|, \\ \Upsilon_2 &:= \max_{(j,k) \in E} n^{-1} \sum_{i=1}^n (\hat{\Theta}_{*j}^T \mathbf{X}_i^{\otimes 2} \hat{\Theta}_{*k} - \hat{\Theta}_{jk} - \Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* + \Theta_{jk}^*)^2. \end{aligned}$$

By Lemma E.4, we are guaranteed that $\sup_{\Theta^* \in \mathcal{M}_{E^n, E^{\text{nc}}} \mathbb{P}(\sqrt{\log|E|}|\Upsilon_1| \geq \sqrt{\log|E|}\xi_1) < \xi_2$, where $\xi_1 = \Xi_1 s \log d / \sqrt{n}$, $\xi_2 = 2/d$ and since under our assumptions $\sqrt{\log|E|}\xi_1 = o(1)$ and $\xi_2 = o(1)$, this satisfies⁴ the first requirement in condition (M) (i) of Theorem 5.1. (Chernozhukov et al., 2013).

Next we turn to bounding $|\Upsilon_2|$. Using (E.7) from Lemma E.6 we conclude that there exists a constant $C > 0$ such that $\sup_{\Theta^* \in \mathcal{M}_{E^n, E^{\text{nc}}} \mathbb{P}((\log(|E|n))^2 \Upsilon_2 \geq Cs^2(\log(|E|n))^2 \log(nd) \log d/n) \leq 2/d$. Clearly since $|E| \leq \binom{d}{2} < d^2$ by our assumption we have $s^2(\log(|E|n))^2 \log(nd) \log d/n = o(1)$. Finally, recall that $\mathbb{E}(\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^* - \Theta_{jk}^*)^2 \geq \delta$ and $\|\Theta_{*j}^{*T} \mathbf{X}_i^{\otimes 2} \Theta_{*k}^*\|_{\psi_1} \leq 2C^4$ as we saw

⁴In fact assumption (M) requires a stronger control but it is not necessary for our purpose.

in the proof of Lemma E.4, and in addition by assumption $(\log(dn))^6/n = o(1)$. This completes the verification of conditions needed by Theorem 5.1 of Chernozhukov et al. (2013) and shows that (3.9) holds.

To see the second part it is sufficient to show that on the first iteration of Algorithm 1, all edges will be rejected. We first obtain a high probability bound on the quantile of $\sqrt{n} \max_{(j,k) \in E} |\tilde{\Theta}_{jk} - \Theta_{jk}^*|$. By Lemma E.4 we get that:

$$\sqrt{n} \max_{(j,k) \in E} |\tilde{\Theta}_{jk} - \Theta_{jk}^*| \leq \sqrt{n} \max_{(j,k) \in E} |\Theta_{*j}^{*T} (\hat{\Sigma} - \Sigma^*) \Theta_{*k}^*| + \Xi_1 s \log d / \sqrt{n},$$

with probability at least $1 - 2/d$, where Ξ_1 is an absolute constant. Next recall that by (E.5) when $\sqrt{\log d/n} \leq 4C^4$ on the same event as above we have:

$$\max_{j,k \in [d]} |\Theta_{*j}^{*T} (\hat{\Sigma} - \Sigma^*) \Theta_{*k}^*| \leq 4\tilde{C}C^4 \sqrt{\log d/n}.$$

Putting last two inequalities together we conclude:

$$\sqrt{n} \max_{(j,k) \in E} |\tilde{\Theta}_{jk} - \Theta_{jk}^*| \leq 4\tilde{C}C^4 \sqrt{\log d} + \Xi_1 s \log d / \sqrt{n}, \quad (\text{E.8})$$

with probability no less than $1 - 2/d$. The last of course implies that any quantile of $\sqrt{n} \max_{(j,k) \in E} |\tilde{\Theta}_{jk} - \Theta_{jk}^*|$ will be smaller than the value on the RHS with high probability. Next, by Corollary 3.1⁵ of Chernozhukov et al. (2013) we have:

$$\sup_{\alpha \in (0,1)} |\mathbb{P}(\sqrt{n} \max_{(j,k) \in E} |\tilde{\Theta}_{jk} - \Theta_{jk}^*| \leq c_{1-\alpha,E}) - (1 - \alpha)| = o(1). \quad (\text{E.9})$$

Combining (E.8) and (E.9) gives us that for large enough d, n , for any fixed $\alpha > 0$ we have:

$$c_{1-\alpha,E} \leq 4\tilde{C}C^4 \sqrt{\log d} + \Xi_1 s \log d / \sqrt{n}.$$

Now, clearly

$$\sqrt{n} |\tilde{\Theta}_{jk}| \geq \sqrt{n} |\kappa \sqrt{\log d/n}| - \sqrt{n} \max_{(j,k) \in E} |\tilde{\Theta}_{jk} - \Theta_{jk}^*| \geq c_{1-\alpha,E}, \quad (\text{E.10})$$

provided that $|\kappa| \geq 9\tilde{C}C^4$, and hence the power goes to 1 uniformly over the parameter space $\mathcal{M}_{E^n, E^{\text{nc}}}$. \square

References

ADDARIO-BERRY, L., BROUTIN, N., DEVROYE, L. and LUGOSI, G. (2010). On combinatorial testing problems. *The Annals of Statistics* **38** pp. 3063–3092.

⁵Formally we verified conditions of Theorem 5.1, but in fact they imply Corollary 3.1 holds.

- ARIAS-CASTRO, E., BUBECK, S. and LUGOSI, G. (2012). Detection of correlations. *The Annals of Statistics* **40** pp. 412–435.
- ARIAS-CASTRO, E., BUBECK, S. and LUGOSI, G. (2015a). Detecting positive correlations in a multivariate sample. *Bernoulli* **21** 209–241.
- ARIAS-CASTRO, E., BUBECK, S., LUGOSI, G. and VERZELEN, N. (2015b). Detecting markov random fields hidden in white noise. *arXiv preprint arXiv:1504.06984* .
- ARIAS-CASTRO, E., CANDÈS, E. J. and DURAND, A. (2011a). Detection of an anomalous cluster in a network. *The Annals of Statistics* **39** pp. 278–304.
- ARIAS-CASTRO, E., CANDÈS, E. J. and PLAN, Y. (2011b). Global testing under sparse alternatives: Anova, multiple comparisons and the higher criticism. *The Annals of Statistics* **39** pp. 2533–2556.
- BARAUD, Y. (2002). Non-asymptotic minimax rates of testing in signal detection. *Bernoulli* **8** pp. 577–606.
- BARTZOKIS, G., BECKSON, M., LU, P. H., NUECHTERLEIN, K. H., EDWARDS, N. and MINTZ, J. (2001). Age-related changes in frontal and temporal lobe volumes in men: a magnetic resonance imaging study. *Archives of General Psychiatry* **58** pp. 461–465.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)* **57** pp. 289–300.
- BERTHET, Q. and RIGOLLET, P. (2013). Optimal detection of sparse principal components in high dimension. *The Annals of Statistics* **41** pp. 1780–1815.
- CAI, T. T., LIU, W. and LUO, X. (2011). A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association* **106** pp. 594–607.
- CAI, W., CHEN, T., SZEGLETES, L., SUPEKAR, K. and MENON, V. (2015). Aberrant cross-brain network interaction in children with attention-deficit/hyperactivity disorder and its relation to attention deficits: a multi-and cross-site replication study. *Biological Psychiatry* .
- CHEN, M., REN, Z., ZHAO, H. and ZHOU, H. (2015). Asymptotically normal and efficient estimation of covariate-adjusted gaussian graphical model. *Journal of the American Statistical Association* **111** pp. 394–406.
- CHERNOZHUKOV, V., CHETVERIKOV, D. and KATO, K. (2013). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics* **41** pp. 2786–2819.
- CHERNOZHUKOV, V., CHETVERIKOV, D. and KATO, K. (2016). Central limit theorems and bootstrap in high dimensions. *Annals of Probability* , to appear.
- DONOHU, D. and JIN, J. (2004). Higher criticism for detecting sparse heterogeneous mixtures. *The Annals of Statistics* **32** pp. 962–994.

- DUBHASHI, D. and RANJAN, D. (1996). Balls and bins: A study in negative dependence. *BRICS Report Series* **3**.
- FANG, K. T., KOTZ, S. and NG, K. W. (1990). *Symmetric multivariate and related distributions*, vol. 36 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, Ltd., London.
- FRIEDMAN, J. H., HASTIE, T. J. and TIBSHIRANI, R. J. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
- GELFAND, A. E., KIM, H.-J., SIRMANS, C. and BANERJEE, S. (2003). Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association* **98** 387–396.
- GU, Q., CAO, Y., NING, Y. and LIU, H. (2015). Local and global inference for high dimensional gaussian copula graphical models. *arXiv preprint arXiv:1502.02347*.
- HELMKE, U. and ROSENTHAL, J. (1995). Eigenvalue inequalities and schubert calculus. *Mathematische Nachrichten* **171** pp. 207–226.
- INGSTER, Y. I. (1982). On the minimax nonparametric detection of signals in white gaussian noise. *Problemy Peredachi Informatsii* **18** pp. 61–73.
- INGSTER, Y. I., TSYBAKOV, A. B. and VERZELEN, N. (2010). Detection boundary in sparse regression. *Electronic Journal of Statistics* **4** pp. 1476–1526.
- JANKOVA, J. and VAN DE GEER, S. A. (2015). Confidence intervals for high-dimensional inverse covariance estimation. *Electronic Journal of Statistics* **9** pp. 1205–1229.
- JOAG-DEV, K. and PROSCHAN, F. (1983). Negative association of random variables with applications. *The Annals of Statistics* **11** pp. 286–295.
- JOHNSTONE, I. M. and LU, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* **104** pp. 682–693.
- LAM, C. and FAN, J. (2009). Sparsistency and rates of convergence in large covariance matrix estimation. *The Annals of Statistics* **37** pp. 4254–4278.
- LE CAM, L. (1973). Convergence of estimates under dimensionality restrictions. *The Annals of Statistics* **1** pp. 38–53.
- LIU, H., HAN, F. and ZHANG, C.-H. (2012). Transelliptical graphical models. *Proc. of NIPS* 809–817.
- LIU, H., LAFFERTY, J. and WASSERMAN, L. (2009). The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *Journal of Machine Learning Research* **10** 2295–2328.
- LIU, W. (2013). Gaussian graphical model estimation with false discovery rate control. *The Annals of Statistics* **41** 2948–2978.
- MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High dimensional graphs and variable selection with the lasso. *The Annals of Statistics* **34** pp. 1436–1462.

- MILHAM, M. P., FAIR, D., MENNES, M., MOSTOFKY, S. H. ET AL. (2012). The ADHD-200 consortium: a model to advance the translational potential of neuroimaging in clinical neuroscience. *Frontiers in systems neuroscience* **6** 62.
- NEYKOV, M., NING, Y., LIU, J. S. and LIU, H. (2015). A unified theory of confidence regions and testing for high dimensional estimating equations. *arXiv preprint arXiv:1510.08986* .
- RASKUTTI, G., YU, B., WAINWRIGHT, M. J. and RAVIKUMAR, P. K. (2008). Model Selection in Gaussian Graphical Models: High-Dimensional Consistency of ℓ_1 -regularized MLE. In *Advances in Neural Information Processing Systems*.
- RAVIKUMAR, P., WAINWRIGHT, M. J., RASKUTTI, G. and YU, B. (2011). High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics* **5** pp. 935–980.
- REN, Z., SUN, T., ZHANG, C.-H. and ZHOU, H. H. (2015). Asymptotic normality and optimalities in estimation of large gaussian graphical models. *The Annals of Statistics* **43** pp. 991–1026.
- ROMANO, J. P. and WOLF, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association* **100** pp. 94–108.
- SANTHANAM, N. P. and WAINWRIGHT, M. J. (2012). Information-Theoretic limits of selecting binary graphical models in high dimensions. *Information Theory, IEEE Transactions on* **58** pp. 4117–4134.
- VERSHYNIN, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing: Theory and Applications* (Y. C. Eldar and G. Kutyniok, eds.). Cambridge University Press.
- VERZELEN, N. and VILLERS, F. (2010). Goodness-of-fit tests for high-dimensional gaussian linear models. *The Annals of Statistics* **38** pp. 704–752.
- WANG, W., WAINWRIGHT, M. J. and RAMCHANDRAN, K. (2010). Information-theoretic bounds on model selection for gaussian markov random fields. In *ISIT*. Citeseer.
- YANG, Y. and BARRON, A. (1999). Information-theoretic determination of minimax rates of convergence. *The Annals of Statistics* **27** pp. 1564–1599.
- YANG, Z., NING, Y. and LIU, H. (2014). On semiparametric exponential family graphical models. *arXiv preprint arXiv:1412.8697* .
- YUAN, M. and LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society. Series B (Methodological)* **68** pp. 49–67.
- ZHANG, C.-H. and ZHANG, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society. Series B (Methodological)* **76** pp. 217–242.